Logical Methods in Computer Science
Vol. 8(3:22)2012, pp. 1–31
www.lmcs-online.org

Submitted Jul. 6, 2011
Published Sep. 21, 2012

TWO FOR THE PRICE OF ONE: LIFTING SEPARATION LOGIC ASSERTIONS

JACOB THAMSBORG [a], LARS BIRKEDAL, AND HONGSEOK YANG [c]

[a,b] IT University of Copenhagen, Rued Langgaards Vej 7, DK-2300 København S, Denmark
e-mail address: {thamsborg, birkedal}@itu.dk

[c] Wolfson Building, Parks Road, Oxford, OX1 3QD, UK
e-mail address: hongseok00@gmail.com

Abstract. Recently, data abstraction has been studied in the context of separation logic, with noticeable practical successes: the developed logics have enabled clean proofs of tricky challenging programs, such as subject-observer patterns, and they have become the basis of efficient verification tools for Java (jStar), C (VeriFast) and Hoare Type Theory (Ynot). In this paper, we give a new semantic analysis of such logic-based approaches using Reynolds's relational parametricity. The core of the analysis is our lifting theorems, which give a sound and complete condition for when a true implication between assertions in the standard interpretation entails that the same implication holds in a relational interpretation. Using these theorems, we provide an algorithm for identifying abstraction-respecting client-side proofs; the proofs ensure that clients cannot distinguish two appropriately-related module implementations.

1. Introduction

Data abstraction is one of the key design principles for building computer software, and it has been the focus of active research from the early days of computer science. Recently, data abstraction has been studied in the context of separation logic [26, 7, 22, 27, 12], with noticeable practical successes: the developed logics have enabled clean proofs of tricky challenging programs, such as the subject-observer pattern, and they have become the basis of efficient verification tools for Java (jStar [14]), C (VeriFast [18]) and Hoare Type Theory (Ynot [23]).

In this paper, we give a new semantic analysis of these logic-based approaches using Reynolds's relational parametricity. Our techniques can be used to prove *representation independence*, i.e., that clients cannot distinguish between related module implementations, a consequence that we would expect from using data abstraction, but (as we shall see) a consequence that only holds for certain good clients.

1998 ACM Subject Classification: F.3.1.
Key words and phrases: separation logic, data abstraction, relational interpretation.

LOGICAL METHODS IN COMPUTER SCIENCE
DOI:10.2168/LMCS-8(3:22)2012
© J. Thamsborg, L. Birkedal, and H. Yang
Creative Commons

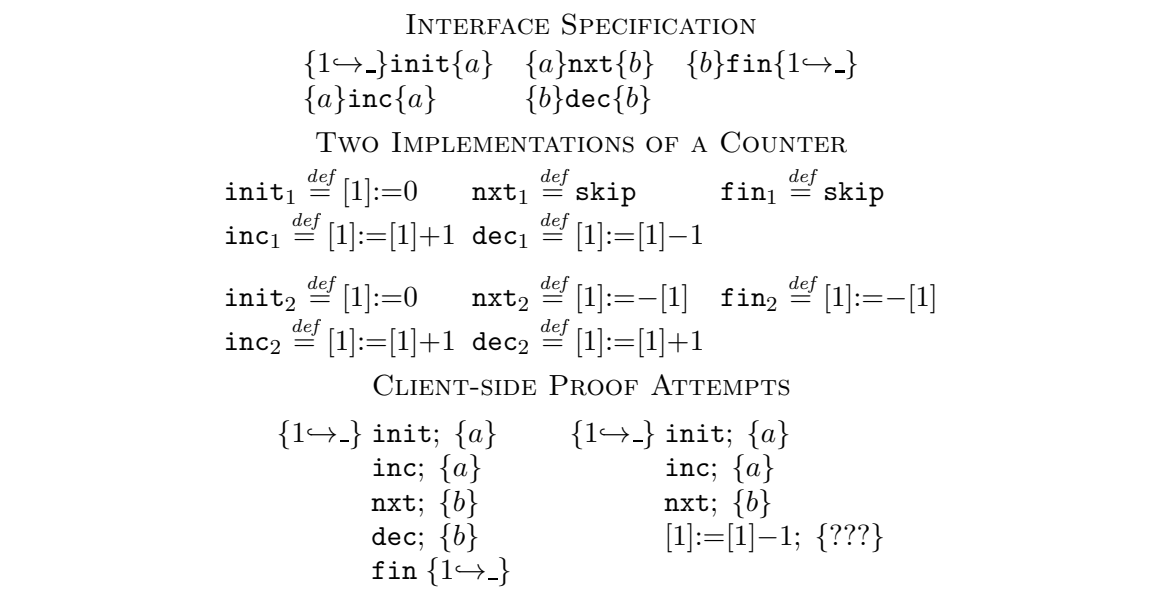


Figure 1: Two-stage Counter

Logic-based Data Abstraction. The basic idea of the logic-based approaches is that the private states of modules are exposed to clients only abstractly using *assertion variables* [7], also known as *abstract predicates* [26][1]. For concreteness, we consider a two-stage counter module and client programs in Figure 1. The module realizes a counter with increment and decrement operations, called `inc` and `dec`. An interesting feature is that the counter goes through two stages in its lifetime; in the first stage, it can perform only the increment operation, but in the second, it can only run the decrement. The interface specification in the figure formalizes this intended behavior of the counter using assertion variables a and b, where a means that the counter is in the first stage and b that the counter is in the second. The triple for `init` says that the initialization can turn the assertion $1\hookrightarrow_$, denoting heaps with cell 1, to the assertion variable a, which describes an abstract state where we can only call `inc` or `nxt` (since a is the precondition of only those operations). The abstract state a can be changed to b by calling `nxt`, says the triple for the `nxt` operation. In b we are allowed to run `dec` but not `inc`. Finally, `fin` can turn the abstract state b back to $1\hookrightarrow_$. Note that by using a and b, the interface specification does not expose the private state of the module to the client. It reveals only *partial information* about the private state of the module; here it is whether the private state is in the first or the second stage. The flexibility afforded by revealing partial information is very useful in applications; see the examples mentioned in the references above.

In these logic-based approaches, proof attempts for clients of a module can succeed only when they are given with respect to the abstract interface specification, without making any further assumptions on assertion variables. For instance, the proof attempt on the bottom left of Figure 1 is successful, whereas the bottom right one is not, because the latter

[1]Abstract predicates do take arguments, though. We conjecture that it is equally expressive to use an assertion variable for each combination of abstract predicate and concrete arguments.

assumes that the assertion variable b entails the allocatedness of cell 1. This is not so, even when the entailment holds for an actual definition of b.

Representation Independence. In this paper, we give a condition on client-side proofs that ensure *representation independence*: take a client with a standard proof of correctness that satisfies this condition and two implementations of a module; if we can relate the heap-usage of the two modules in a way preserved by the module operations, then the client gives the same result with both modules. To relate the heap-usage, we need to give, for each assertion variable, a relation on heaps and verify that the module operations respect these relations — the *coupling relations*.

As an example, consider the left-hand side client in Figure 1. The proof of the specification

$$\{1\hookrightarrow_\}\ \texttt{init};\ \texttt{inc};\ \texttt{nxt};\ \texttt{dec};\ \texttt{fin}\ \{1\hookrightarrow_\}$$

satisfies the forthcoming condition on client-side proofs. Also, we have two implementations of the counter module that the client makes use of; both use cell 1 to represent their private states, but in different ways — the first stores the current value of the counter, but the second stores the current value or its negative version, depending on whether it is in the first stage or the second. Accordingly, we give the two coupling relations:

$$r_a \stackrel{def}{=} \{(h_1,h_2) \mid 1\in \mathrm{dom}(h_1)\cap \mathrm{dom}(h_2) \wedge h_1(1)=h_2(1)\}$$
$$r_b \stackrel{def}{=} \{(h_1,h_2) \mid 1\in \mathrm{dom}(h_1)\cap \mathrm{dom}(h_2) \wedge h_1(1)=-h_2(1)\}$$

It is easy to see that all module operations preserve these coupling relations. If, say, $(h_1,h_2) \in r_b$, then we have $h_1(1) = n$ and $h_2(1) = -n$ for some n and so $(h_1[1 \mapsto n-1], h_2[1 \mapsto -n+1]) \in r_b$ too; hence the decrement operations of the modules respect the coupling relation. By Theorem 6.1 we now get that the client specification is valid also in a binary reading: if we take any two heaps and run the client with one module in the first and with the other module in the second then we will end up with two heaps holding the same value in cell 1 — provided that we started out with two such. Indeed, the binary reading of the assertion $1\hookrightarrow_$ is

$$\{(h_1,h_2) \mid 1\in \mathrm{dom}(h_1)\cap \mathrm{dom}(h_2) \wedge h_1(1) = h_2(1)\}$$

which incidentally coincides with r_a.

It is worthwhile to emphasize that this is not a consequence of the standard unary reading of the specification of the client: due to the existentially quantified content of cell 1, running with the one module could yield different contents of cell 1 than running with the other module, even if the contents are initially the same. On the other hand, it is the presence of this quantification that makes the binary reading worthwhile: if our client had a more exhaustive specification, say

$$\{1\hookrightarrow 0\}\ \texttt{init};\ \texttt{inc};\ \texttt{nxt};\ \texttt{dec};\ \texttt{fin}\ \{1\hookrightarrow 0\},$$

then the standard unary reading suffices for representation independence and the binary reading would provide no news. Often, though, the more exhaustive specification will be harder to prove, in particular for verification tools.

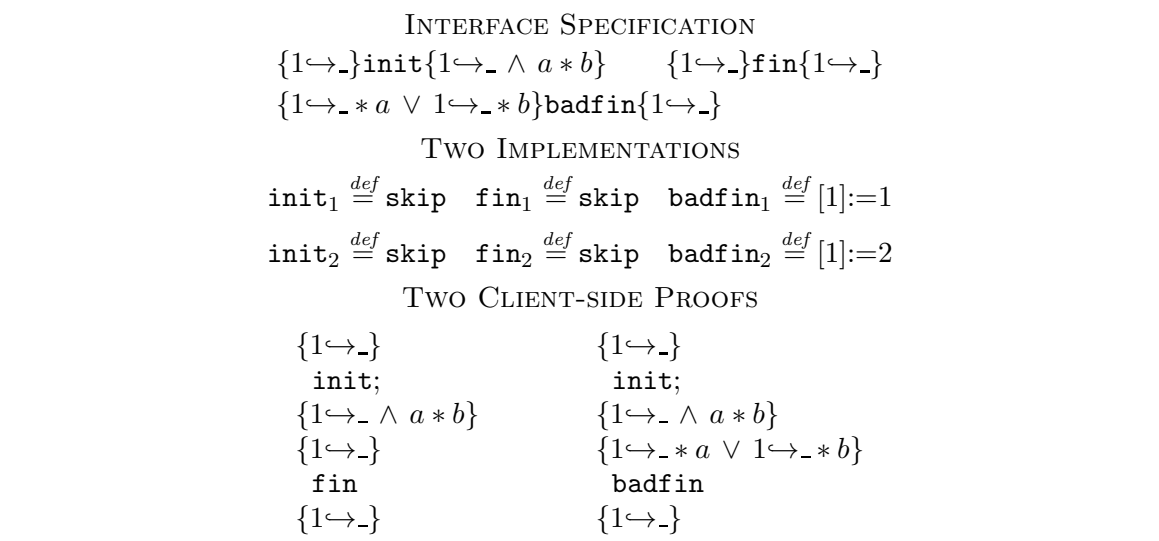


Figure 2: Good or Bad Client-side Proofs

The Rule of Consequence and Lifting. In earlier work [10] we were able to prove such a representation independence result for a more restricted form of logical data abstraction, namely one given by frame rules rather than general assertion variables. Roughly speaking, frame rules use a restricted form of assertion variables that are not exposed to clients at all, as can be seen from some models of separation logic in which frame rules are modelled via quantification over semantic assertions [9]. This means that the rules do not allow the exposure of even partial information about module internals. (On the other hand, frame rules implement information hiding, because they completely relieve clients of tracking the private state of a module, even in an abstracted form.) Our model in [10] exploited this restricted use of assertion variables, and gave relational meanings to Hoare triple specifications, which led to representation independence.

Removing this restriction and allowing assertion variables in client proofs turned out to be very challenging. The challenge is the use of the rule of consequence in client-side proofs; this has implications between assertions (possibly containing assertion variables) as hypotheses, and such do not always *lift*, i.e., they may hold in the standard, unary reading of assertion whilst failing in the binary reading. In this paper, we provide a sound and, in a certain sense, complete answer to when the lifting can be done.

For instance, consider the example in Figure 2. Our results let us conclude that the client-side proof on the left is *good* but the one on the right is *bad*; hence we expect to derive representation independence only from the former. The client on the left calls `init` and ends with the post-condition $(1\hookrightarrow_ \wedge a * b)$. Since $(1\hookrightarrow_ \wedge a * b) \implies 1\hookrightarrow_$ is true in the standard interpretation, the rule of consequence can be applied to yield the precondition of `fin`, which can be called, ending up with the postcondition $(1\hookrightarrow_)$. The key point here is the implication used in the rule of consequence. Our results imply that this implication can indeed be lifted to an implication between relational meanings of assertions $(1\hookrightarrow_ \wedge a * b)$ and $1\hookrightarrow_$ (Theorem 4.5 in Section 4). They also entail that this lifting implies the representation

independence theorem. The coupling relations

$$r_a \stackrel{def}{=} \{(h_1, h_2) \mid 1 \in \mathrm{dom}(h_1)\}$$
$$r_b \stackrel{def}{=} \{(h_1, h_2) \mid 1 \in \mathrm{dom}(h_2)\}$$

are preserved by the modules: the relational meaning of $1 \hookrightarrow _ * a \vee 1 \hookrightarrow _ * b$ is empty; note that separating conjunctions binds more tightly than conjunction and disjunction. Hence the client on the left should give the same result for both modules and, indeed, both $\mathtt{init}_1; \mathtt{fin}_1$ and $\mathtt{init}_2; \mathtt{fin}_2$ are the same skip command.

The client on the right also first calls `init` and then uses the rule of consequence. But this time our results say that a true implication $(1 \hookrightarrow _ \wedge a * b) \implies (1 \hookrightarrow _ * a \vee 1 \hookrightarrow _ * b)$ in the rule of consequence does *not* lift to an implication between relational meanings of the assertions: the pair of heaps $([1 \mapsto 0], [1 \mapsto 0])$ belong to the left hand side but possibly not to the right if the pair $([], [1 \mapsto 0])$ is in the interpretation of a and the pair $([1 \mapsto 0], [])$ is in the interpretation of b; see Example 4.9 for details. Because of this failure, the proof of the client does not ensure representation independence. In fact, the client can indeed distinguish between the two module implementations — when the client is executed with the first module implementation, the final heap maps address 1 to 1, but when the client is executed with the second, the final heap maps address 1 to 2.

Note that we phrase the lifting only in terms of semantically true implications, without referring to how they are proved. By doing so, we make our results relevant to automatic tools that use the semantic model of separation logic to prove implications, such as the ones based on shallow embeddings of the assertion logic [23, 16].

To sum up, the question of whether representation independence holds for a client with a proof comes down to whether, in the proof, the implications used in the rule of consequence can be lifted to a relational interpretation. In this paper, we give a sound and, in a certain sense, complete characterization of when that holds.

It is appropriate to remark already here, that although we extend our assertions with assertion variables we also restrict them to contain neither ordinary nor separating implication. And, in the end, we consider only a fragment of those. Details are given in Sections 2 and 4; here we just remark that the assertions we do study are not unlike the ones considered in the tools mentioned in the beginning of this introduction, in particular jStar. Also we use intuitionistic separation logic as we envision a language with garbage collection; this, too, is in line with the jStar tool.

The rest of the paper is organized as follows:

Sections 2 and 3: give the meanings of assertions, both the standard and relational meanings. Indeed, we give, for any $n \in \mathsf{PosInt}$, the n-ary meaning of assertions as n-ary relations on the set of heaps. These relations are intuitionistic, i.e., they are upward closed relations with respect to heap extension.

Section 4: contain the main technical contributions of the paper. We give, for assertions of a particular form, a sound and, in a certain sense, complete answer to the question of when we may lift implications between assertions from the standard, unary meaning to the binary meaning.

Section 5: has the curious spinoff result that an implication between assertions holds for arbitrary arity if and only if it holds for reasons of parametric polymorphism in a particular sense.

Section 6: returns to the main line of development; this is where we show that a client-side proof yields representation independence if it uses the rule of consequence only with implications that lift.

Section 7: concludes the paper.

Proofs are found in the appendices, the main text only give details for a few examples.

2. Semantic Domain

In the following section we will define the meaning of an assertion to be an n-ary relation on heaps. To formalize this relational meaning, we need a semantic domain IRel_n of relations, which we define and explain in this section.

Let Heap be the set of finite partial functions from positive integers to integers (i.e., $\mathsf{Heap} \stackrel{def}{=} \mathsf{PosInt} \rightarrow_{fin} \mathsf{Int}$), ranged over by f, g, h. This is a commonly used set for modelling heaps in separation logic, and it has a partial commutative monoid structure $([], \cdot)$, where $[]$ is the empty heap and the $\cdot$ operator combines disjoint heaps:

$$[] \stackrel{def}{=} \emptyset, \qquad f \cdot g \stackrel{def}{=} \begin{cases} f \cup g & \text{if } \mathrm{dom}(f) \cap \mathrm{dom}(g) = \emptyset \\ \text{undefined} & \text{otherwise} \end{cases}$$

The operator $\cdot$ induces a partial order $\sqsubseteq$ on Heap, modelling heap extension, by $f \sqsubseteq g$ iff $g = f \cdot h$ for some h.

We also consider the $+$ operator for combining possibly-overlapping but consistent heaps, and the $-$ operator for subtracting one heap from another:

$$f + g \stackrel{def}{=} \begin{cases} f \cup g & \text{if } \forall l \in \mathrm{dom}(f) \cap \mathrm{dom}(g).\ f(l) = g(l) \\ \text{undefined} & \text{otherwise} \end{cases}$$

$$(f - g)(l) \stackrel{def}{=} \begin{cases} f(l) & \text{if } l \in \mathrm{dom}(f) \setminus \mathrm{dom}(g) \\ \text{undefined} & \text{otherwise} \end{cases}$$

We call an n-ary relation $r \subseteq \mathsf{Heap}^n$ *upward closed* iff $(f_1, \ldots, f_n) \in r \wedge (\forall i.\ f_i \sqsubseteq g_i) \implies (g_1, \ldots, g_n) \in r$.

Definition 2.1. IRel_n is the family of upward closed n-ary relations on heaps.

Note that IRel_1 consists of upward closed *sets* of heaps, which are frequently used to interpret assertions in separation logic for garbage-collected languages. We call elements of IRel_1 *predicates* and denote them by p, q.

For every $n \geq 1$, domain IRel_n has a complete lattice structure: join and meet are given by union and intersection, bottom is the empty relation, and top is Heap^n. The domain also has a semantic separating conjunction connective defined by

$$\begin{aligned}(f_1, .., f_n) \in r * s \stackrel{def}{\iff} \exists (g_1, .., g_n) \in r.\ \exists (h_1, .., h_n) \in s.& \\ (g_1, .., g_n) \cdot (h_1, .., h_n) = (f_1, .., f_n).&\end{aligned}$$

Here we use the component-wise extension of $\cdot$ for tuples. Intuitively, a tuple is related by $r * s$ when it can be split into two disjoint tuples, one related by r and the other by s.

The domain IRel_1 of predicates is related to IRel_n for every n, by the map $\Delta_n \stackrel{def}{=} \lambda p.\{(f, \ldots, f) \mid f \in p\}^{\uparrow}$, where $\uparrow$ is the upward closure on relations. Note that each predicate is turned into an n-ary identity relation on p modulo the upward closure. This map behaves well with respect to the structures discussed on IRel_1 and IRel_n, as expressed by the lemma below:

Lemma 2.2. *Function Δ_n preserves the complete lattice structure and the $*$ operator.*

For every $n \geq 1$, the domain IRel_n has further structure: it has standard, semantic separating implication and upwards-closed implication; as such, it is a complete BI algebra [8]. Unfortunately, the above lemma fails for both implications. And that lemma is the pivot of the upcoming results; it is the basic link between the unary and binary (and n-ary) readings of assertions. This is why we leave out these connectives in our assertions in the next section; it is a fundamental limitation in our approach.

3. Assertions and Relational Semantics

Let Var and AVar be disjoint sets of normal variables $x, y, ...$ and assertion variables $a, b, ...$, respectively. Our assertions φ are from separation logic, and they conform to the following grammar:

$$E ::= x \mid 0 \mid 1 \mid E + E \mid \dots \qquad P ::= E \hookrightarrow E \mid \dots$$
$$\varphi ::= P \mid a \mid \varphi * \varphi \mid \mathsf{true} \mid \varphi \wedge \varphi \mid \mathsf{false} \mid \varphi \vee \varphi$$
$$\mid \forall x.\, \varphi \mid \exists x.\, \varphi$$

In the grammar, E is a heap-independent expression, and P is a primitive predicate, which in the standard interpretation denotes an upward closed set of heaps. For instance, $E \hookrightarrow E'$ means heaps containing cell E with contents E'. The dots in the grammar indicate possible extensions of cases, such as multiplication for E and inductive predicates for P. We will use the abbreviation $E \hookrightarrow _$ for $\exists y. E \hookrightarrow y$.

An assertion φ is given a meaning $[\![\varphi]\!]^n_{\eta,\rho} \in \mathsf{IRel}_n$ as an n-ary relation on heaps, where the arity n is a parameter of the interpretation. Here environment η maps normal variables in φ to integers, and ρ maps assertion variables in φ to n-ary relations in IRel_n. When φ does not contain any assertion variables, we often omit ρ and write $[\![\varphi]\!]^n_\eta$, because the meaning of φ does not depend on ρ. We will make use of unary and binary semantics most places, but in Section 5 we will explore higher arities as well.

We define the semantics of φ, using the complete lattice structure and the $*$ operator of the domain IRel_n; see Figure 3. Note that the relational semantics of primitive predicates is defined by embedding their standard meanings via Δ_n. In fact, this embedding relationship holds for all assertions without assertion variables, because Δ_n preserves the semantic structures of the domains (Lemma 2.2):

Lemma 3.1. *For all φ and η, ρ, ρ', if $\Delta_n(\rho(a)) = \rho'(a)$ for every $a \in \mathsf{AVar}$, we have that $\Delta_n([\![\varphi]\!]^1_{\eta,\rho}) = [\![\varphi]\!]^n_{\eta,\rho'}$.*

We write $\varphi \models^n \psi$ to mean that $[\![\varphi]\!]^n_{\eta,\rho} \subseteq [\![\psi]\!]^n_{\eta,\rho}$ holds for all environments η, ρ. If n=1, this reduces to the standard semantics of assertions in separation logic. We will use the phrase "$\varphi \Longrightarrow \psi$ is n-ary valid" to mean that $\varphi \models^n \psi$ holds. In addition, we write $\varphi \models^n_\eta \psi$ for a fixed η to mean that $[\![\varphi]\!]^n_{\eta,\rho} \subseteq [\![\psi]\!]^n_{\eta,\rho}$ holds for all environments ρ; we say that "$\varphi \implies \psi$ is n-ary η-valid" if this is true.

4. Lifting Theorems and Completeness

We call an assertion φ *simple* if it is of the form $(\bigvee_{i=1}^{I} \bigwedge_{j=1}^{J} \varphi_{(i,j)} * \mathbf{a}_{(i,j)})$, where $\mathbf{a}_{(i,j)}$ is a vector of assertion variables and $\varphi_{i,j}$ is an assertion not containing any assertion variables. We will consider the question of lifting an implication between simple assertions φ, ψ to a binary relational interpretation: when does $\varphi \models^1 \psi$ imply that $\varphi \models^2 \psi$?

$$
\begin{array}{ll}
[\![P]\!]^n_{\eta,\rho} \stackrel{def}{=} \Delta_n((\!|P|\!)\eta) & [\![\varphi * \psi]\!]^n_{\eta,\rho} \stackrel{def}{=} [\![\varphi]\!]^n_{\eta,\rho} * [\![\psi]\!]^n_{\eta,\rho} \\
[\![a]\!]^n_{\eta,\rho} \stackrel{def}{=} \rho(a) & [\![\varphi \wedge \psi]\!]^n_{\eta,\rho} \stackrel{def}{=} [\![\varphi]\!]^n_{\eta,\rho} \cap [\![\psi]\!]^n_{\eta,\rho} \\
[\![\mathsf{true}]\!]^n_{\eta,\rho} \stackrel{def}{=} \mathsf{Heap}^n & [\![\varphi \vee \psi]\!]^n_{\eta,\rho} \stackrel{def}{=} [\![\varphi]\!]^n_{\eta,\rho} \cup [\![\psi]\!]^n_{\eta,\rho} \\
[\![\mathsf{false}]\!]^n_{\eta,\rho} \stackrel{def}{=} \emptyset & [\![\forall x.\varphi]\!]^n_{\eta,\rho} \stackrel{def}{=} \bigcap_{v \in \mathsf{Int}} [\![\varphi]\!]^n_{\eta[x \mapsto v],\rho} \\
 & [\![\exists x.\varphi]\!]^n_{\eta,\rho} \stackrel{def}{=} \bigcup_{v \in \mathsf{Int}} [\![\varphi]\!]^n_{\eta[x \mapsto v],\rho}
\end{array}
$$

where $(\!|P|\!)\eta$ is the standard semantics of P as an upward closed set of heaps, which satisfies:

$$(\!|E \hookrightarrow F|\!)\eta \;=\; \{f \mid [\![E]\!]\eta \in \mathrm{dom}(f) \;\wedge\; f([\![E]\!]\eta) = [\![F]\!]\eta\}.$$

Figure 3: Interpretation of Assertions

The simple assertions are a fragment of the assertions considered in the above section: simple assertions are not, in general, closed under separating conjunction as the latter does not distribute over conjunction, nor are quantified simple assertions necessarily simple. The divide, however, between simple and non-simple assertions is not deep. The forthcoming completeness result is intimately connected to the form of the assertions, but it is very possible that the basic ideas from lifting could be applied to a larger fragment. We have not considered that to any extent, however. It is worth mentioning that all assertions considered in Section 1 are simple. On the other hand, for assertions φ_1, φ_2 and φ_3 with no assertion variables and assertion variables a_1,a_2 and a_3, we do not, in general, have simplicity of an assertion like

$$\big[(\varphi_1 * a_1) \wedge (\varphi_2 * a_2)\big] * \big[\varphi_3 * a_3\big].$$

It should be noted, that simple assertions include most of the important aspects of the fragments of separation logic used by automatic program analysis tools. For instance, if we ignore so called primed variables (which correspond to existentially-quantified variables), the original SpaceInvader uses separation-logic formulas of the form $\bigvee_{i=1}^{I}(P_{i,1} * \ldots * P_{i,k_i})$ [13], and its most recent extension for handling a particular class of graph-like data structures uses $\bigwedge_{j=i}^{J}\bigvee_{i=1}^{I}(P_{i,j,1} * \ldots * P_{i,j,k_{i,j}})$ [19]. Note that in both cases, either formulas are already simple or they can be easily transformed to equivalent simple formulas. The assertions used by the jStar tool [15] has neither ordinary implication, separating implication nor ordinary conjunction and only quite restricted use of quantifiers. Since proofs obtained from such tools are one target of our results, we argue that the restrictions imposed on assertions are not unreasonable in terms of usage.

The starting point of our analysis is to realize that it is sufficient to study implications of the form:

$$\bigwedge_{i=1}^{M} \varphi_i * a_{i,1} * \cdots * a_{i,M_i} \implies \bigvee_{j=1}^{N} \psi_j * b_{j,1} * \cdots * b_{j,N_j} \tag{4.1}$$

where φ_i's and ψ_j's do not contain assertion variables, and no assertion variables occur only on the right hand side of the implication.

Lemma 4.1. *There is an algorithm taking simple assertions φ, ψ and returning finitely many implications $\{\varphi^l \implies \psi^l\}_{l \in L}$, such that (a) $\varphi^l \implies \psi^l$ has the form (4.1) for all $l \in L$, and (b) for any $n \in \{1, 2\}$, we have that $\varphi \models^n \psi$ holds iff $\varphi^l \models^n \psi^l$ holds for all $l \in L$.*

The algorithm in the lemma is given in Appendix B.

Thus, in this section, we will focus on lifting implications of the form (4.1). Specifically, we will give a *complete* answer to the following question: Given one such implication that is η-valid in the unary interpretation for some environment η, can we decide if the implication is η-valid in the binary interpretation *merely* by inspection of the layout of the assertion variables? The answer will come in two parts. The first part, in Section 4.3, provides three lifting theorems, each of which has a criterion on the variable layout that, if met, implies that η-validity may be lifted regardless of the φ_i's and ψ_j's. The second part, in Section 4.4, is a completeness theorem; it states that if the variables fail the criteria of all three lifting theorems then there are choices of φ_i's and ψ_j's with no variables such that we have unary but not binary validity.

This approach has pros and cons. Assume that we have an implication of the aforementioned form that is valid in the unary interpretation, and we would like to know if it is valid in the binary interpretation too. Trying out the layout of the variables against the criteria of the three lifting theorems is an easily decidable and purely syntactical process – and if it succeeds then we have binary validity. If it fails, however, we are at a loss; we know that there are φ_i's and ψ_j's with the same variable layout such that lifting fails but we do not learn anything about our concrete implication. There is, however, an alternate use of the theory below if the lifting criteria fail; we will elaborate on that in Section 6.

4.1. **Notation.** We need some notation that will accompany us throughout this section. Consider an implication of the form (4.1). Let $V = \bigcup_{i=1}^{M} \{a_{i,1}, \ldots, a_{i,M_i}\}$ be the set of all left hand side assertion variables, these include the right hand side assertion variables too by assumption. Define $\Pi : \{1, \ldots, M\} \to \mathsf{Nat}^V$ and $\Omega : \{1, \ldots, N\} \to \mathsf{Nat}^V$ by the following:

$$\Pi(i)(c) \stackrel{def}{=} |\{k \mid a_{i,k} \equiv c\}|, \quad \Omega(j)(c) \stackrel{def}{=} |\{k \mid b_{j,k} \equiv c\}|.$$

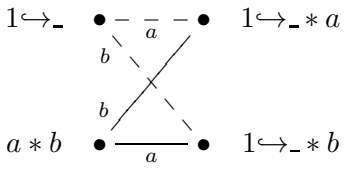
These functions give vectors of assertion variable counts for each conjunct and disjunct. For $1 \leq i \leq M$ and $1 \leq j \leq N$ we write $\Pi(i) \geq \Omega(j)$ if we have $\Pi(i)(c) \geq \Omega(j)(c)$ for each variable $c \in V$, i.e., if conjunct i has the same or a greater number of occurrences of all variables than disjunct j. We write $\Pi(i) \not\geq \Omega(j)$ if this fails, i.e., if there is $c \in V$ such that $\Pi(i)(c) < \Omega(j)(c)$. If a conjunct, say conjunct i, has no variables, i.e., if $\Pi(i)(c) = 0$ holds for all $c \in V$, then we say it is *empty*; the same goes for the disjuncts.

We shall write $-$ to denote $\exists n, m.\ n{\hookrightarrow}m$, meaning heaps with at least one cell. On the semantic side, we write $[m]$ for $m \in \mathsf{PosInt}$ to denote the heap that stores 0 at location m and nothing else. For $m_0, ..., m_n \in \mathsf{PosInt}$ different we write $[m_0, ..., m_n]$ for $[m_0] \cdot ... \cdot [m_n]$.

Finally we introduce a piece of sanity-preserving graphical notation. We depict an implication of the form (4.1) as a complete bipartite graph with the conjuncts lined up on the left hand side and the disjuncts on the right hand side. For any $1 \leq i \leq M$ and any $1 \leq j \leq N$ we draw a solid line from conjunct i to disjunct j if $\Pi(i) \geq \Omega(j)$. We label that line with all the $c \in V$ such that $\Pi(i)(c) > \Omega(j)(c)$ if indeed there are any such. If, on the other hand, $\Pi(i) \not\geq \Omega(j)$ then we draw a dashed line instead and label it with all the $c \in V$ such that $\Pi(i)(c) < \Omega(j)(c)$. Note that the drawing of edges depend solely on the layout of the variables; the φ_i's and ψ_j's have no say in the matter. As an example, the implication

$$1{\hookrightarrow}_ \wedge a * b \implies 1{\hookrightarrow}_ * a \vee 1{\hookrightarrow}_ * b,$$

which we shall look into in Example 4.9, is depicted as follows:



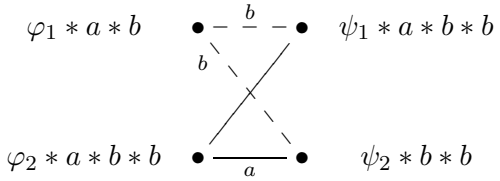
With a little experience, it is quite easy to check the conditions of the upcoming lifting theorems by looking at the corresponding graph; the graphs expose the structure of the assertions important to the proofs.

4.2. **Strategy.** We give a brief strategic overview before the onslaught. Consider an implication of the form (4.1). If the layout of the variables satisfy (at least) one of three upcoming criteria then we know this: unary η-validity holds only if it holds for 'obvious reasons'. The latter is captured precisely in the Parametricity Condition but, loosely, it says that there are $1 \leq i \leq M$ and $1 \leq j \leq N$ such that $\phi_i \implies \psi_j$ is η-valid in the unary interpretation and such that $\Pi(i) \geq \Omega(j)$. This is sufficiently parametric in the treatment of assertion variables that it immediately implies binary η-validity and even n-ary η-validity for any n.

The three criteria, as given in the next subsection, are rather technical; each is what it takes for proof idea of the corresponding lifting theorem to go through. They are complete, however: if the implication fails all three criteria then there are choices of φ_i's and ψ_j's such that unary η-validity holds for 'non-obvious reasons'; in particular such that binary η-validity fails. Non-obvious reasons comes down to exploiting the limited arity in different ways; we elaborate on that in Subsection 4.4.

4.3. **Layouts that Lift.** The following is a first example of a layout of variables that ensure that for any choice of φ_i's and ψ_j's we get that unary η-validity of the implication yields binary η-validity. That it holds is a consequence of Theorem 4.5 but we have spelled out a concrete proof that will serve as a guide to the further development.

Example 4.2 (Shadow-Lift)**.** For any four assertions $\varphi_1, \varphi_2, \psi_1, \psi_2$ with no assertion variables and any appropriate environment η we have that unary η-validity of the following implication implies binary η-validity:



Assume that we have unary η-validity. Before we go on to consider the binary case we derive a simple unary consequence that does not involve assertion variables: For any $h \in \mathsf{Heap}$ with subheaps $h_1 \sqsubseteq h$ and $h_2 \sqsubseteq h$ such that $h_1 \in [\![\varphi_1]\!]^1_\eta$ and $h_2 \in [\![\varphi_2]\!]^1_\eta$ we get that $h_2 \in [\![\psi_1]\!]^1_\eta$ or that $h_2 \in [\![\psi_2]\!]^1_\eta$.

To prove this, let h, h_1 and h_2 be as assumed. We build $\rho : \{a, b\} \to \mathsf{IRel}_1$ by letting $\rho(a) = \mathsf{Heap}$ and letting $\rho(b)$ be the following union of sets of heaps:

$$\{(h - h_1) \cdot [n, n+1]\}^\uparrow \cup \{(h - h_2) \cdot [n]\}^\uparrow \cup \{[n+1]\}^\uparrow$$

where $n = \max(\mathrm{dom}(h) \cup \{0\}) + 1$. It is now immediate that $h \cdot [n, n+1]$ lies in the interpretation of both conjuncts since

$$h_1 \cdot (h - h_1) \cdot [n, n+1] = h \cdot [n, n+1] = h_2 \cdot (h - h_2) \cdot [n] \cdot [n+1],$$

and so by our assumption on the original implication it must lie in the interpretation on one of the disjuncts too. Suppose that we have

$$h \cdot [n, n+1] \in [\![\psi_1 * a * b * b]\!]^1_{\eta,\rho} = [\![\psi_1]\!]^1_\eta * \rho(b) * \rho(b),$$

where the equality holds because $\rho(a) = \mathsf{Heap}$ is the unit for $*$. We then write $h \cdot [n, n+1] = g_1 \cdot g_2 \cdot g_3$ for $g_1 \in [\![\psi]\!]^1_\eta$ and $g_2, g_3 \in \rho(b)$. But as g_2 and g_3 have disjoint domains we must have $(h - h_2) \cdot [n] \sqsubseteq g_2$ and $[n+1] \sqsubseteq g_3$ or the version with g_2 and g_3 swapped. In any case we have that

$$\begin{aligned} \mathrm{dom}(g_1) &= \mathrm{dom}(h \cdot [n, n{+}1]) \setminus (\mathrm{dom}(g_2 \cdot g_3)) \\ &\subseteq \mathrm{dom}(h \cdot [n, n{+}1]) \setminus (\mathrm{dom}(h{-}h_2) \cup \{n, n{+}1\}) \\ &= \mathrm{dom}(h_2). \end{aligned}$$

But then we have $g_1 \sqsubseteq h_2$ and since $g_1 \in [\![\psi_1]\!]^1_\eta$ we get $h_2 \in [\![\psi_1]\!]^1_\eta$ too. If we have $h \cdot [n, n+1] \in [\![\psi_2 * b * b]\!]^1_{\eta,\rho}$ we proceed similarly.

The above short proof is the crux of the example. It implies unary η-validity – this we knew already – but also the binary η-validity. To see this, we pick an arbitrary environment $\rho : \{a, b\} \to \mathsf{IRel}_2$, we take arbitrary $(h_1, h_2) \in [\![\varphi_1 * a * b \wedge \varphi_2 * a * b * b]\!]^2_{\eta,\rho}$ and we aim to prove that $(h_1, h_2) \in [\![\psi_1 * a * b * b \vee \psi_2 * b * b]\!]^2_{\eta,\rho}$ too. We split (h_1, h_2) according to the conjuncts. Because of Lemma 3.1 and the upward closedness condition of IRel_2, we can write

$$(h_1, h_2) = (g^1, g^1) \cdot (g^2_1, g^2_2) \cdot (g^3_1, g^3_2)$$

for $g^1 \in [\![\varphi_1]\!]^1_\eta$, $(g^2_1, g^2_2) \in \rho(a)$ and $(g^3_1, g^3_2) \in \rho(b)$. Also we can write

$$(h_1, h_2) = (f^1, f^1) \cdot (f^2_1, f^2_2) \cdot (f^3_1, f^3_2) \cdot (f^4_1, f^4_2)$$

for $f^1 \in [\![\varphi_2]\!]^1_\eta$, $(f^2_1, f^2_2) \in \rho(a)$ and $(f^3_1, f^3_2), (f^4_1, f^4_2) \in \rho(b)$. But now $g^1 + f^1$ with subheaps g^1 and f^1 meet the conditions of the unary consequence from above, and so we get $f^1 \in [\![\psi_1]\!]^1_\eta$ or $f^1 \in [\![\psi_2]\!]^1_\eta$ and the second splitting of (h_1, h_2) shows that (h_1, h_2) lie in the binary interpretation of the first or second disjunct, respectively. Notice that neither $g^1 \in [\![\psi_1]\!]^1_\eta$ nor $g^1 \in [\![\psi_2]\!]^1_\eta$ would have worked since the first conjunct has too few variables, i.e., $\Pi(1) \not\geq \Omega(1)$ and $\Pi(1) \not\geq \Omega(2)$ □

The simple idea justifies the odd choice of name: we attach to each occurrence of b in the conjuncts a 'shadow' in such a way that any two shadows from different conjuncts overlap. This means that the two occurrences of b in, say, the first disjunction must correspond to occurrences of b in the same conjunct; in particular that conjunct must have at least two occurrences. We attach no shadow to a, instead we remove a by instantiating to Heap; this is because the second disjunct lacks an occurrence of a and hence we may fail to 'peel off' the entire shadow. Had a occurred as the single label of a dashed line, this removal would have 'introduced' a solid line and the approach would fail.

Generalizing the unary consequence that served as the crucial stepping stone in the above example we arrive at the following condition on our implications:

Definition 4.3 (Parametricity Condition)**.** Assume that we have an implication of the form (4.1) and an appropriate environment η. We say that the Parametricity Condition is

satisfied if, for all $h, h_1, \ldots, h_M \in \mathsf{Heap}$ with $h_i \sqsubseteq h$ and $h_i \in [\![\varphi_i]\!]^1_\eta$ for all $1 \leq i \leq M$, it is is the case that one (or both) of the following conditions hold:

(1) There are $1 \leq i \leq M$ and $1 \leq j \leq N$ such that $h_i \in [\![\psi_j]\!]^1_\eta$ and $\Pi(i) \geq \Omega(j)$.
(2) There is $1 \leq j \leq N$ such that $h \in [\![\psi_j]\!]^1_\eta$ and the j-th disjunct is empty.

Note that specializing the Parametricity Condition, henceforth just the PC, to an implication of the form treated in the above example yields the stated unary consequence because no disjuncts are empty. The second option in the PC will be motivated later.

We emphasize that the PC may hold or may fail for any given combination of an implication and environment η. But if it holds then we have binary η-validity; the proof in case of the first option of the PC is an easy generalization of the latter half of the above example:

Proposition 4.4. *The PC implies binary η-validity.*

We arrive now at the first lifting theorem. It is a generalization of the former half of Example 4.2; the proof of the theorem has a lot more details to it than the example but the overall idea is the same. The theorem states a criterion on the layout of the variables that, if met, means that unary η-validity implies the PC and hence also binary η-validity. The criterion is, loosely, that we can remove all variables that occur as labels of solid lines without introducing new solid lines and without emptying any disjuncts:

Theorem 4.5 (Shadow-Lift)**.** *Unary η-validity of an implication implies the PC if each dashed line is labeled with at least one variable which is not a label on a solid line and each disjunct has an occurrence of a variable that is not a label on a solid line. Spelling it out in symbols, we require, with $L = \{(i,j) \mid 1 \leq i \leq M \wedge 1 \leq j \leq N\}$, that*

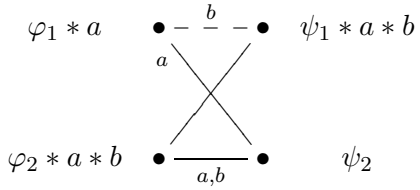
$$\begin{array}{l} \forall (i,j) \in L.\ \ \Pi(i) \not\geq \Omega(j) \implies \\ \quad \exists c \in V.\ \ \Pi(i)(c) < \Omega(j)(c) \ \wedge \\ \qquad \big(\forall (k,l) \in L.\ \ \Pi(k) \geq \Omega(l) \implies \Pi(k)(c) = \Omega(l)(c)\big) \end{array}$$

and

$$\begin{array}{l} \forall 1 \leq j \leq N.\ \ \exists c \in V.\ \ \Omega(j)(c) > 0 \ \wedge \\ \quad \big(\forall (k,l) \in L.\ \ \Pi(k) \geq \Omega(l) \implies \Pi(k)(c) = \Omega(l)(c)\big). \end{array}$$

As motivation for the next lifting theorem, we note that the variable layout criterion of the above theorem fails if one or more disjuncts are empty. Correspondingly, we never touch upon the second option of the PC. But there are variable layouts with empty disjuncts that ensure lifting:

Example 4.6 (Balloon-Lift)**.** For any four assertions $\varphi_1, \varphi_2, \psi_1, \psi_2$ with no assertion variables and any appropriate environment η we have that unary η-validity of the following implication implies binary η-validity:



Assume unary η-validity. As in Example 4.2 we derive a unary consequence as an intermediate result: For any $h \in \mathsf{Heap}$ with subheaps $h_1 \sqsubseteq h$ and $h_2 \sqsubseteq h$ such that $h_1 \in [\![\varphi_1]\!]^1_\eta$ and $h_2 \in [\![\varphi_2]\!]^1_\eta$ we have that either $h_2 \in [\![\psi_1]\!]^1_\eta$ or $h \in [\![\psi_2]\!]^1_\eta$.

To prove this, let h, h_1 and h_2 be as assumed. We construct $\rho : \{a, b\} \to \mathsf{IRel}_1$ by letting $\rho(a) = \mathsf{Heap}$ and $\rho(b) = \{h - h_2\}^\uparrow$. We get that

$$h \sqsupseteq h_1 \in [\![\varphi_1]\!]^1_\eta = [\![\varphi_1]\!]^1_\eta * \mathsf{Heap} = [\![\varphi_1 * a]\!]^1_{\eta,\rho},$$

and

$$h = h_2 \cdot (h - h_2) \in [\![\varphi_2]\!]^1_\eta * \rho(b) = [\![\varphi_2]\!]^1_\eta * \mathsf{Heap} * \rho(b) = [\![\varphi_2 * a * b]\!]^1_{\eta,\rho}.$$

This means that h must lie in the interpretation of one of the disjuncts. If it is the first, we inspect the interpretation and get that

$$h = g_1 \cdot g_2 \cdot g_3$$

for $g_1 \in [\![\psi_1]\!]^1_\eta$, $g_2 \in \mathsf{Heap}$ and $g_3 \sqsupseteq h - h_2$. It means that

$$\begin{aligned} \mathrm{dom}(g_1) &= \mathrm{dom}(h) \setminus \mathrm{dom}(g_2 \cdot g_3) \subseteq \mathrm{dom}(h) \setminus \mathrm{dom}(g_3) \\ &\subseteq \mathrm{dom}(h) \setminus \mathrm{dom}(h - h_2) = \mathrm{dom}(h_2) \end{aligned}$$

which implies that $g_1 \sqsubseteq h_2$ and so $h_2 \in [\![\psi_1]\!]^1_\eta$. If, on the other hand, h lies in the interpretation of the second disjunct then we are done immediately.

Now we prove the claim of binary η-validity. We pick an arbitrary environment $\rho : \{a, b\} \to \mathsf{IRel}_2$, we take arbitrary $(h_1, h_2) \in [\![\varphi_1 * a \wedge \varphi_2 * a * b]\!]^2_{\eta,\rho}$ and we must prove that $(h_1, h_2) \in [\![\psi_1 * a * b \vee \psi_2]\!]^2_{\eta,\rho}$ too. We write

$$(h_1, h_2) = (g^1, g^1) \cdot (g^2_1, g^2_2)$$

for $g^1 \in [\![\varphi_1]\!]^1_\eta$ and $(g^2_1, g^2_2) \in \rho(a)$, and

$$(h_1, h_2) = (f^1, f^1) \cdot (f^2_1, f^2_2) \cdot (f^3_1, f^3_2)$$

for $f^1 \in [\![\varphi_2]\!]^1_\eta$, $(f^2_1, f^2_2) \in \rho(a)$ and $(f^3_1, f^3_2) \in \rho(b)$. But now $g^1 + f^1$ with subheaps g^1 and f^1 satisfies the above properties and so we get $f^1 \in [\![\psi_1]\!]^1_\eta$ or $g^1 + f^1 \in [\![\psi_2]\!]^1_\eta$. If $f^1 \in [\![\psi_1]\!]^1_\eta$ holds then the second splitting of (h_1, h_2) shows that (h_1, h_2) is in the interpretation of the first disjunct. If $g^1 + f^1 \in [\![\psi_2]\!]^1_\eta$, we are done too, since we may write $(h_1, h_2) = (g^1 + f^1, g^1 + f^1) \cdot (e_1, e_2)$ for some $(e_1, e_2) \in \mathsf{Heap}^2$ and so (h_1, h_2) lies in the interpretation $[\![\psi_2]\!]^2_\eta = \Delta([\![\psi_2]\!]^1_\eta)$ of the second conjunct. □

Once again, the underlying idea is simple: we attach 'shadows' to occurrences of variables, but this time we stay within the the original heap. This is quite inhibitory as we may have to use the empty heap as shadow. Again we remove a variable (in general a set of variables) by instantiating to Heap but the remaining variable (in general the remaining set of variables) must satisfy quite restrictive demands.

Just as we did for Example 4.2 we may generalize the former half of this example yielding Theorem 4.7 below. The latter half of the example, on the other hand, constitutes an example of the approach of the proof of Proposition 4.4 in case we run into the second option of the PC. Note also that specializing the PC to an implication of the form considered in the example yields the stated unary consequence.

Theorem 4.7 (Balloon-Lift)**.** *Unary η-validity of an implication implies the PC if there is a subset $B \subseteq V$ with the following three properties. First, each conjunct has at most one occurrence of a variable from B, i.e.,*

$$\forall 1 \leq i \leq M.\ \sum_{c \in B} \Pi(i)(c) \leq 1.$$

Second, each disjunct is empty or has exactly one occurrence of a variable from B, i.e.,

$$\forall 1 \leq j \leq N.\ \sum_{c \in V} \Omega(j)(c) = 0 \ \vee \ \sum_{c \in B} \Omega(j)(c) = 1.$$

Third, each dashed line must have a label from B. That is, when $L = \{(i,j) \mid 1 \leq i \leq M \wedge 1 \leq j \leq N\}$,

$$\forall (i,j) \in L.\ \Pi(i) \not\geq \Omega(j) \implies \exists c \in B.\ \Pi(i)(c) < \Omega(j)(c).$$

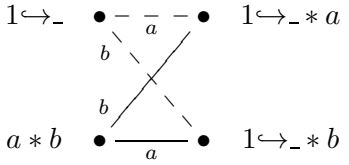
One thing to note about the theorem is that if we have no empty disjuncts, none of the variables in the subset $B \subseteq V$ can be labels of a solid line. In particular, the conditions of Theorem 4.5 are met, so the above theorem is really only useful if one or more disjuncts are empty. A simple but pleasing observation is that this theorem is applicable if all conjuncts and all disjuncts have at most a single occurrence of any assertion variable; in that case, we can just choose $B = V$ above.

The final lifting theorem captures the oddities of the special case of just one conjunct:

Theorem 4.8 (Lonely-Lift)**.** *Unary η-validity of an implication implies the PC if there is just one conjunct, i.e., M=1, and all lines are solid, i.e., $\Pi(1) \geq \Omega(j)$ for all $1 \leq j \leq N$.*

4.4. **Completeness.** It is now time for examples of implications that do not lift, i.e., that are valid in the unary interpretation but not in the binary. The first is based on the following observation: If $h \in [\![1 \hookrightarrow _]\!]^1_\eta$ and $h \in p * q$ for $h \in \mathsf{Heap}$ and $p, q \in \mathsf{IRel}_1$ then we have $h \in [\![1 \hookrightarrow _]\!]^1_\eta * p$ or $h \in [\![1 \hookrightarrow _]\!]^1_\eta * q$. This is because we must have $[1 \mapsto n] \sqsubseteq h$ for some $n \in \mathsf{Int}$ and so writing $h = h_1 \cdot h_2$ with $h_1 \in p$ and $h_2 \in q$ gives us $[1 \mapsto n] \sqsubseteq h_1$ or $[1 \mapsto n] \sqsubseteq h_2$. But this line of argument breaks down if we change to binary reading. We have, e.g., $([1],[1]) \in [\![1 \hookrightarrow _]\!]^2_\eta$ and $([1],[1]) \in \{([1],[])\}^\uparrow * \{([],[1])\}^\uparrow$ but both $[\![1 \hookrightarrow _]\!]^2_\eta * \{([1],[])\}^\uparrow$ and $[\![1 \hookrightarrow _]\!]^2_\eta * \{([],[1])\}^\uparrow$ are empty. We can recast this as an implication that cannot be lifted:

Example 4.9 (Fan-Counter)**.** This implication is valid on the unary but not on the binary level:



First we argue that the implication holds on the unary level. Let $\rho : \{a, b\} \to \mathsf{IRel}_1$ be an arbitrary environment of upwards closed sets of heaps to a and b. Let $h \in \mathsf{Heap}$ be arbitrary and assume that

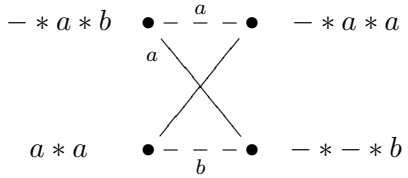
$$h \ \in \ [\![1 \hookrightarrow _ \wedge (a * b)]\!]^1_{\eta,\rho} \ = \ [\![1 \hookrightarrow _]\!]^1_\rho \cap (\rho(a) * \rho(b)).$$

By the above observation we get either $h \in [\![1 \hookrightarrow _]\!]^1_\eta * \rho(a)$ or $h \in [\![1 \hookrightarrow _]\!]^1_\eta * \rho(b)$ which matches the right hand side of the implication.

Now we move on to prove that the implication fails on the binary level. Define an environment $\rho : \{a, b\} \to \mathsf{IRel}_2$ by $\rho(a) = \{([1],[])\}^\uparrow$ and $\rho(b) = \{([],[1])\}^\uparrow$. Then, $[\![1 \hookrightarrow _ \wedge a * b]\!]^2_{\eta,\rho} = [\![1 \hookrightarrow _]\!]^2_{\eta,\rho} \cap (\rho(a) * \rho(b))$, which contains the pair $([1],[1])$. But, as observed, both disjuncts have empty binary interpretations. □

An observation of similar nature is that for $p \in \mathsf{IRel}_1$ we have either $p = \mathsf{Heap}$ or $p \subseteq [\![-]\!]^1_\eta = \{[m{\mapsto}n] \mid m \in \mathsf{PosInt}, n \in \mathsf{Int}\}^\uparrow$ because if $p \neq \mathsf{Heap}$ then it cannot contain the empty heap. On the binary level, however, we have $\mathsf{Heap}^2 \neq \{([1],[])\}^\uparrow \not\subseteq [\![-]\!]^2_\eta = \{([m{\mapsto}n],[m{\mapsto}n]) \mid m \in \mathsf{PosInt}, n \in \mathsf{Int}\}^\uparrow$. One consequence is this:

Example 4.10 (Bridge-Counter)**.** This implication is valid on the unary but not on the binary level:



First we argue that the implication holds on the unary level. Let $\rho : \{a, b\} \to \mathsf{IRel}_1$ be an arbitrary environment that assigns upwards closed sets of heaps to each of the two variables. We branch on the value of $\rho(a)$. If $\rho(a) \neq \mathsf{Heap}$ then we have $\rho(a) \subseteq [\![-]\!]^1_\eta$ which again means that the first conjunct directly implies the second disjunct. If $\rho(a) = \mathsf{Heap}$ holds, we get that

$$\begin{aligned}[\![- * a * b]\!]^1_{\eta,\rho} &= [\![-]\!]^1_{\eta,\rho} * \mathsf{Heap} * \rho(b) = [\![-]\!]^1_{\eta,\rho} * \rho(b) \\ &\subseteq [\![-]\!]^1_{\eta,\rho} = [\![-]\!]^1_{\eta,\rho} * \mathsf{Heap} * \mathsf{Heap} = [\![- * a * a]\!]^1_{\eta,\rho}\end{aligned}$$

because Heap is the unit for $*$. Hence we get that the first conjunct implies the first disjunct and we have proved that the implication holds unarily.

Now we prove that the implication fails on the binary level. Define an environment $\rho : \{a, b\} \to \mathsf{IRel}_2$ by $\rho(a) = \{([1],[])\}^\uparrow \cup \{([2],[2])\}^\uparrow$ and $\rho(b) = \mathsf{Heap}^2$. Observe now that $([1,2],[2]) = ([2],[2]) \cdot ([1],[]) \cdot ([],[])$, which implies that $([1,2],[2]) \in [\![- * a * b]\!]^2_{\eta,\rho}$. From the rewriting $([1,2],[2]) = ([1],[]) \cdot ([2],[2])$, we get $([1,2],[2]) \in [\![a * a]\!]^2_{\eta,\rho}$ too and so this pair of heaps lies in the interpretation of the left hand side. But it does not belong to the interpretation of either disjunct. An easy – if somewhat indirect – way of realizing this is to note that any pair of heaps in either $[\![-]\!]^2_{\eta,\rho}$ or in $[\![a * a]\!]^2_{\eta,\rho}$ must have a second component with nonempty domain. But then any pair of heaps in the interpretation of either disjunct must have a second component with a domain of at least two elements. In particular, neither can contain the pair $([1,2],[2])$. □

In principle, the above two observations are all that we need to prove completeness. Or, phrased differently, assume that we have a layout of variables that fail the criteria of all three lifting theorems; by applying one of the two observations, we can then build a concrete implication with that variable layout and with unary but not binary validity.

Having said that, the territory to cover is huge; the full completeness proof is a lengthy and rather technical journey, the details of which do not provide much insight. We supply it as a series of lemmas in Appendix D; these include generalizations of Example 4.9 and Example 4.10 above. If one verifies the lemmas in the order listed and apply them as sketched then it is feasible, if not exactly easy, to prove the following:

Theorem 4.11 (Completeness)**.** *If a variable layout meets none of the criteria in Theorems 4.5, 4.7 and 4.8, then there are choices of φ_i's and ψ_j's with no variables such we have unary but not binary validity.*

4.5. **Future Work: Supported Assertions.** By now, we have given a complete division of the possible layouts of variables into those that lift and those that do not. The divide is technical; without some understanding of the underlying proofs, it is hard to get an intuitive feel for it.

One way to simplify would be to consider *supported* assertions. A n-ary relation $r \in \mathsf{IRel}_n$ is supported if, for every $\mathbf{f} \in \mathsf{Heap}^n$, it holds that $\mathbf{g_1} \sqsubseteq \mathbf{f}$ and $\mathbf{g_2} \sqsubseteq \mathbf{f}$ and $\mathbf{g_1}, \mathbf{g_2} \in r$ together implies the existence of $\mathbf{g} \in r$ with $\mathbf{g} \sqsubseteq \mathbf{g_1}$ and $\mathbf{g} \sqsubseteq \mathbf{g_2}$. In an intuitionistic setting, the supported assertions play the role that the precise assertions do in a classical, non-intuitionistic setting: they validate reasoning about the resource invariant in the Hypothetical Frame Rule [25] and about shared resources in concurrent separation logic [11]. The problems we face here appear reminiscent and a natural question is this: how about restricting assertion variables to supported assertions?

We have not investigated this in any detail, but initial findings suggest that this would simplify matters, maybe extensively so. The counter-example given in Example 4.9 still holds, so it is not the case that everything lifts, but the counter-example of Example 4.10 breaks. The central proof of Theorem 4.5 uses non-supported assertions; on the other hand, if $r \in \mathsf{IRel}_n$ is supported then either $r * r$ is empty or we have $r = \mathsf{Heap}^n$ so maybe we could restrict to conjuncts with at most one occurrence of each assertion variable.

Along the same lines, it would be interesting to revisit the challenges in a classical, non-intuitionistic setting. This, too, is left for future work with the one comment that the counter-example given in Example 4.9 persists.

5. Higher Arities and Parametricity

We saw in Proposition 4.4 that the PC implies binary η-validity of an implication. It is easy to show that the PC also implies unary η-validity, either directly or by observing that binary implies unary. A natural question to ask is whether we can reverse this. Example 4.9 shows that unary validity does not entail the PC, because the latter fails for that concrete implication. But as binary validity fails too, we could hope that binary validity would enforce the PC. Unfortunately, this is not the case, as demonstrated by the implication

$$1 \hookrightarrow _ \wedge a * a * b \implies 1 \hookrightarrow _ * a \vee 1 \hookrightarrow _ * b.$$

Here the PC is the same as for Example 4.9 and hence still is not true, but we do have binary validity. We do not, however, have ternary validity but the example could easily be scaled: having n occurrences of a in the second conjunct means n-ary but not $n+1$-ary validity for any $n \geq 1$. In summary, we have seen that for any $n \geq 1$ we can have n-ary validity whilst the PC fails.

What does hold, however, is the following:

Theorem 5.1. *For an implication of the form* (4.1) *and an appropriate environment η we have that n-ary η- validity implies the PC if $n \geq \max\{2, M_1, \ldots, M_M\}$.*

Notice how this fits nicely with the above example: with n occurrences of a we have n-ary validity but we need $(n+1)$-ary validity to prove the PC since there is also a single b. The proof is in Appendix E, and reuses techniques from the proofs of Theorems 4.5 and 4.7.

By an easy generalization of Proposition 4.4 we have the following corollary to the above theorem:

Corollary 5.2. *The PC holds iff we have n-ary η-validity for all $n \geq 1$.*

This corollary can be read, loosely, as a coincidence between parametric polymorphism as introduced by Strachey [31] and relational parametricity as proposed by Reynolds [30]: The PC corresponds to Strachey parametricity in the loose sense that if it holds, then there is an approach, parametric in the assertion variables, that produce right hand side proofs of heap membership from the left hand side ones: Take a heap, split it along the conjuncts, apply the PC to the parts in the interpretations of the φ's and you are done, possibly after discarding some variables. This involves no branching or other intrinsic operations on the assertion variables, which we are free to discard by our intuitionistic setup. If, on the other hand, the implication is η-valid for arbitrary arity, then it is fair to call it relationally parametric. Note also that the Examples 4.9 and 4.10 branch on assertion variable values.

This result is analogous to the conjecture of coincidence between Strachey parametricity and n-ary relational parametricity for traditional type-based parametricity [29, Page 2].

Finally we note that as a consequence of the above corollary we have that the lifting theorems in the previous section really show that unary validity can be lifted to validity of arbitrary arity. In some sense, they are stronger than required for representation independence, for which binary validity suffices. The authors are unaware of any practical applications of this fact.

6. Representation Independence

In this section, we relate our lifting theorems to representation independence. We consider separation logic with assertion variables where the rule of consequence is restricted according to our lifting theorems, and we define a relational semantics of the logic, which gives a representation independence theorem: all proved clients cannot distinguish between appropriately related module implementations.

To keep the presentation simple, we omit while-loops and allocation from the language. Adding the former together with the standard proof rule is straightforward. Allocation, however, is non-trivial: the notion of having *one* client using *two* modules will be hard-coded into our relational reading of the logic, and allocation on part of the client must give the same address when run with either module. This fails with standard, non-deterministic allocation; it was resolved earlier, however, by Birkedal and Yang [10] using a combination of FM sets and continuations and that approach is applicable here too.

We consider commands C given by the grammar:

$$C \ ::= \ k \mid [E]{:=}E \mid \mathbf{let}\, y{=}[E]\, \mathbf{in}\, C \mid C; C \mid \mathbf{if}\, B\, C\, C$$

Here B is a heap-independent boolean expression, such as $x{=}0$. Commands C are from the loop-free simple imperative language. They can call module operations k, and manipulate heap cells; command $[x]{:=}E$ assigns E to the heap cell x, and this assigned value is read by $\mathbf{let}\, y{=}[x]\, \mathbf{in}\, C$, which also binds y to the read value and runs C under this binding.

Properties of commands C are specified using Hoare triples $\Gamma \vdash \{\varphi\}C\{\psi\}$, where the context Γ is a set of triples for module operations. Figure 4 shows rules for proving these properties. In the figure, we omit contexts, if the same context Γ is used for all the triples.

The rule of consequence deserves attention. Note that the rule uses semantic implications $\models^1$ in the standard unary interpretation, thus allowing the use of existing theorem provers for separation logic. The rule does not allow all semantic implications, but only those that pass our algorithm Chk, so as to ensure that the implications can lift to the

$$\frac{\textsc{Chk}(\varphi',\varphi) \quad \varphi' \models^1 \varphi \quad \{\varphi\}C\{\psi\} \quad \psi \models^1 \psi' \quad \textsc{Chk}(\psi,\psi')}{\{\varphi'\}C\{\psi'\}}$$

$$\frac{\{\varphi\}C\{\varphi'\}}{\{\varphi * \psi\}C\{\varphi' * \psi\}} \qquad \frac{\{\varphi\}C\{\psi\}}{\{\exists x.\, \varphi\}C\{\exists x.\, \psi\}}\ x \notin \mathsf{FV}(C)$$

$$\frac{}{\Gamma, \{\varphi\}k\{\psi\} \vdash \{\varphi\}k\{\psi\}} \qquad \frac{}{\{E \hookrightarrow _\}[E] := F\{E \hookrightarrow F\}}$$

$$\frac{\{\varphi * E \hookrightarrow x\}C\{\psi\}}{\{\exists x. \varphi * E \hookrightarrow x\}\mathbf{let}\ x = [E]\ \mathbf{in}\ C\{\psi\}}\ x \notin \mathsf{FV}(\psi)$$

$$\frac{\{\varphi\}C\{\varphi'\} \quad \{\varphi'\}C'\{\psi\}}{\{\varphi\}C; C'\{\psi\}} \qquad \frac{\{\varphi \wedge B\}C\{\psi\} \quad \{\varphi \wedge \neg B\}C'\{\psi\}}{\{\varphi\}\mathbf{if}\ B\ C\ C'\{\psi\}}$$

Figure 4: Proof Rules

relational level. Our algorithm $\textsc{Chk}(\varphi, \psi)$ performs two checks, and returns true only when both succeed. The first check is whether φ and ψ can be transformed to simple assertions φ' and ψ', using only the distribution of $*$ over $\exists x$ and $\vee$ and distributive lattice laws for $\vee$ and $\wedge$. If this check succeeds and gives φ' and ψ', the algorithm transforms $\varphi' \models^1 \psi'$ to a set of implications of the form (4.1) in Section 4 (Lemma 4.1). Then, for each implication in the resulting set, it tests if any of the the three criteria for lifting are met and returns true if that is always the case[2].

Commands C are interpreted in a standard way, as functions of the type: $[\![C]\!]_{\eta,u} \in \mathsf{Heap} \to (\mathsf{Heap} \cup \{err\})$. Here err denotes a memory error, and η and u are environments that provide the meanings of, respectively, free ordinary variables and module operations. For instance, $[\![k]\!]_{\eta,u}$ is $u(k)$.

Our semantics of triples, on the other hand, is not standard, and uses the binary interpretation of assertions: $(\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}C\{\psi\}$ iff

$$\forall r \in \mathsf{IRel}_2.\ \forall f, g \in \mathsf{Heap}.\ \ (f,g) \in [\![\varphi]\!]^2_{\eta,\rho} * r \implies (([\![C]\!]_{\eta,u_1}(f), [\![C]\!]_{\eta,u_2}(g)) \in [\![\psi]\!]^2_{\eta,\rho} * r.$$

The environment ρ provides the meanings of assertion variables, and the 2-dimensional vector $\mathbf{u}$ gives the two meanings for module operations; intuitively, each u_i corresponds to the i-th module implementation. The interpretation means that if two module implementations $\mathbf{u}$ are used by the same client C, then these combinations should result in the same computation, in the sense that they map φ-related input heaps to ψ-related outputs. The satisfaction of triples can be extended to $(\eta, \rho, \mathbf{u}) \models^2 \Gamma$, by asking that all triples in Γ should hold wrt. $(\eta, \rho, \mathbf{u})$. Using these satisfaction relations on triples and contexts, we define the notion of 2-validity of judgements: $\Gamma \vdash \{\varphi\}C\{\psi\}$ is 2-valid iff

$$\forall(\eta, \rho, \mathbf{u}).\ (\eta, \rho, \mathbf{u}) \models^2 \Gamma \implies (\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}C\{\psi\}.$$

Theorem 6.1. *Every derivable $\Gamma \vdash \{\varphi\}C\{\psi\}$ is 2-valid.*

[2]Recall that the failure of the lifting theorems do not imply that a concrete implication cannot be lifted; consider, e.g., Example 4.9 and replace $1 \hookrightarrow _$ with true everywhere. One can sidestep the general lifting theorems and (try to) verify directly the Parametricity Condition from Definition 4.3 for all environments η. It is, however, a semantic condition and probably undecidable in general.

It is this theorem that we use to derive the representation independence results mentioned in the introduction. Consider again the example in Figure 1. Since the proof of the left hand side client C is derivable using the above rules in the context

$$\Gamma = \{1\hookrightarrow_\}\texttt{init}\{a\}, \{a\}\texttt{inc}\{a\}, \{a\}\texttt{nxt}\{b\},$$
$$\{b\}\texttt{dec}\{b\}, \{b\}\texttt{fin}\{1\hookrightarrow_\},$$

we get 2-validity of $\Gamma \vdash \{1\hookrightarrow_\}C\{1\hookrightarrow_\}$. Instantiating, in the definition of 2-validity, ρ with the given coupling relations and $\mathbf{u}$ with the module implementations gives us

$$(\emptyset, \rho, \mathbf{u}) \models^2 \{1\hookrightarrow_\}C\{1\hookrightarrow_\},$$

since we already know that the different operations respect the coupling relations. Therefore, when we run the client C with the related module implementations, we find that C maps $[\![1\hookrightarrow_]\!]^2$-related heaps (i.e, heaps with the same value at cell 1) to $[\![1\hookrightarrow_]\!]^2$-related heaps again.

7. Conclusion and Discussion

In this paper, we have given a sound and, in a certain sense, complete characterization of when *semantic* implications in separation logic with assertion variables can be lifted to a relational interpretation. This characterization has, then, been used to identify proofs of clients that respect the abstraction of module internals, specified by means of assertion variables, and to show representation independence for clients with such proofs. We hope that our results provide a solid semantic basis for recent logic-based approaches to data abstraction.

In earlier work, Banerjee and Naumann [2] studied relational parametricity for dynamically allocated heap objects in a Java-like language. Later they extended their results to cover clients programs that are correct with respect to specifications following the "Boogie methodology" implemented in the Spec# verifier [3, 4]. In both works, Banerjee and Naumann made use of a non-trivial semantic notion of confinement to describe internal resources of a module; here instead we use separation logic with assertions variables to describe which resources are internal to the module.

Data abstraction and information hiding have been studied in logics and specification languages for pointer programs, other than separation logic. Good example projects are ESC-Modular-3 [20], ESC-Java [17] and Spec# [5], some of which use concepts analogous to abstract predicates, called abstract variables [21]. It would be an interesting future direction to revisit the questions raised in the paper in the context of these logics and specification languages.

Relational interpretations have also been used to give models of programming languages with local state, which can validate representation independence results [24, 28, 6, 1]. These results typically rely on the module allocating the private state, whereas we use the power of separation logic and allow the ownership transfer of states from client to module. For instance, in the two-stage counter in the introduction, the ownership of the cell 1 is transferred from the client to the module upon calling `init`. Even with this ownership transfer, representation independence is guaranteed, because we consider only those clients having (good) proofs in separation logic. This contrasts with representation independence results in local state models, which consider not some but all well-typed clients. The work by Banerjee and Naumann [2] discussed above also permits ownership transfer.

Acknowledgement

We would like to thank the anonymous referees for thorough reviews and insightful comments; their work made this work better. We are also grateful to John Wickerson and Peter O'Hearn for their helpful comments. Yang was supported by EPSRC.

References

[1] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *POPL*, pages 340–353, 2009.
[2] A. Banerjee and D. Naumann. Ownership confinement ensures representation independence for object-oriented programs. *JACM*, 52(6), 2005.
[3] A. Banerjee and D. A. Naumann. State based ownership, reentrance, and encapsulation. In *ECOOP*, pages 387–411, 2005.
[4] M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *FMCO*, pages 364–387, 2005.
[5] M. Barnett, R. DeLine, M. Fähndrich, B. Jacobs, K. R. M. Leino, W. Schulte, and H. Venter. The spec# programming system: Challenges and directions. In *VSTTE*, pages 144–152, 2005.
[6] N. Benton and B. Leperchey. Relational reasoning in a nominal semantics for storage. In *TLCA*, pages 86–101, 2005.
[7] B. Biering, L. Birkedal, and N. Torp-Smith. BI hyperdoctrines and higher-order separation logic. In *ESOP*, 2005.
[8] B. Biering, L. Birkedal, and N. Torp-Smith. BI-hyperdoctrines, higher-order separation logic, and abstraction. *TOPLAS*, 29(5), 2007.
[9] L. Birkedal, N. Torp-Smith, and H. Yang. Semantics of separation-logic typing and higher-order frame rules. In *LICS*, pages 260–269, 2005.
[10] L. Birkedal and H. Yang. Relational parametricity and separation logic. In *FOSSACS*, pages 93–107, 2007.
[11] S. D. Brookes. A semantics for concurrent separation logic. In P. Gardner and N. Yoshida, editors, *CONCUR*, volume 3170 of *Lecture Notes in Computer Science*, pages 16–34. Springer, 2004.
[12] W.-N. Chin, C. David, H. H. Nguyen, and S. Qin. Enhancing modular oo verification with separation logic. In *POPL*, 2008.
[13] D. Distefano, P. W. O'Hearn, and H. Yang. A local shape analysis based on separation logic. In *TACAS*, pages 287–302, 2006.
[14] D. Distefano and M. Parkinson. jStar: towards practical verification for java. In *OOPSLA*, pages 213–226, 2008.
[15] D. Distefano and M. Parkinson J. jstar: Towards practical verification for java. *ACM Sigplan Notices*, 43(10):213–226, 2008.
[16] R. Dockins, A. Hobor, and A. Appel. A fresh look at separation algebras and share accounting. In *APLAS*, pages 161–177, 2009.
[17] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for java. In *PLDI*, 2002.
[18] B. Jacobs and F. Piessens. The VeriFast program verifier. Technical Report CW-520, Katholieke Universiteit Leuven, 2008.
[19] O. Lee, H. Yang, and R. Petersen. Program analysis for overlaid data structures. In *CAV*, 2011.
[20] K. R. M. Leino and G. Nelson. An extended static checker for modular-3. In *CC*, pages 302–305, 1998.
[21] K. R. M. Leino and G. Nelson. Data abstraction and information hiding. *ACM Trans. Program. Lang. Syst.*, 24(5):491–553, 2002.
[22] A. Nanevski, A. Ahmed, G. Morrisett, and L. Birkedal. Abstract predicates and mutable ADTs in Hoare type theory. In *ESOP*, 2007.
[23] A. Nanevski, G. Morrisett, A. Shinnar, P. Govereau, and L. Birkedal. Ynot: dependent types for imperative programs. In *ICFP*, pages 229–240, 2008.
[24] P. O'Hearn and R. Tennent. Parametricity and local variables. *JACM*, 42(3):658–709, May 1995.
[25] P. W. O'Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. *ACM Trans. Program. Lang. Syst.*, 31(3), 2009.

[26] M. Parkinson and G. Bierman. Separation logic and abstraction. In *POPL*, pages 247–258, 2005.
[27] M. Parkinson and G. Bierman. Separation logic, abstraction and inheritance. In *POPL*, pages 75–86, 2008.
[28] A. Pitts and I. Stark. Observable properties of higher order functions that dynamically create local names, or: What's new? In *MFCS*, pages 122–141, 1993.
[29] G. Plotkin and M. Abadi. A logic for parametric polymorphism. In *TLCA*, pages 361–375, 1993.
[30] J. C. Reynolds. Types, abstraction, and parametric polymorphism. In *Information Processing 83*, pages 513–523, 1983.
[31] C. Strachey. Fundamental concepts in programming languages. In *Proceedings of the 1967 International Summer School in Computer Programming, Copenhagen, Denmark.*, 1967.

Appendix A. Proofs of Lemma 2.2 and Lemma 3.1

Lemma 2.2. *Function Δ_n preserves the complete lattice structure and the $*$ operator.*

Proof. From the definition, it is immediate that $\Delta_n(\mathsf{Heap}) = \mathsf{Heap}^n$ and $\Delta_n(\emptyset) = \emptyset$, the former because we have $[] \in \mathsf{Heap}$. Now consider a non-empty family $\{p_i\}_{i\in I}$ of predicates in IRel_1. In order to show the preservation of the complete lattice structure, we need to prove that

$$\Delta_n(\bigcap_{i\in I} p_i) = \bigcap_{i\in I} \Delta_n(p_i) \quad \wedge \quad \bigcup_{i\in I} \Delta_n(p_i) = \Delta_n(\bigcup_{i\in I} p_i).$$

The $\subseteq$ direction in both cases is easy; it follows from the monotonicity of Δ_n.

We start with the $\supseteq$ direction for the meet operator. Pick $(h_1, \ldots, h_n)$ from $\bigcap_{i\in I} \Delta_n(p_i)$. Then,

$$\forall i \in I.\ (h_1, \ldots, h_n) \in \Delta_n(p_i).$$

By the definition of Δ_n, this means that

$$\forall i \in I.\ \exists f_i \in p_i.\ f_i \sqsubseteq h_1 \ \wedge \ \ldots \ \wedge \ f_i \sqsubseteq h_n. \tag{A.1}$$

Let $f = \sum_{i\in I} f_i$. The sum here is well-defined, because (a) there are only finitely many f's such that $f \sqsubseteq h_k$ for all $1 \le k \le n$, and (b) any two such f and g should have the same value for every location in $\mathrm{dom}(f) \cap \mathrm{dom}(g)$. Since all f_i's satisfy (A.1), their sum f also satisfies

$$f \sqsubseteq h_1 \ \wedge \ \ldots \ \wedge \ f \sqsubseteq h_n.$$

Furthermore, $f \in \bigcap_{i\in I} p_i$, because p_i's are upward closed and f is an extension of f_i in p_i. Hence, $\Delta_n(\bigcap_{i\in I} p_i) \subseteq \bigcap_{i\in I} \Delta_n(p_i)$.

Next we prove the $\supseteq$ direction for the join operator. Pick $(h_1, \ldots, h_n)$ from $\Delta_n(\bigcup_{i\in I} p_i)$. Then,

$$\exists i \in I.\ \exists f \in p_i.\ f \sqsubseteq h_1 \ \wedge \ \ldots \ \wedge \ f \sqsubseteq h_n.$$

Hence, by the definition of Δ_n,

$$(h_1, \ldots, h_n) \ \in \ \Delta_n(p_i) \ \subseteq \ \bigcup_{i\in I} \Delta_n(p_i),$$

as desired.

Finally, it remains to show that Δ_n preserves the $*$ operator. Consider predicates $p, q \in \mathsf{IRel}_1$. We need to prove that

$$\Delta_n(p * q) \ = \ \Delta_n(p) * \Delta_n(q).$$

Choose an arbitrary $(h_1, \ldots, h_n)$ from $\Delta_n(p * q)$. By the definition of $\Delta_n(p * q)$, it follows that

$$\exists f \in p.\ \exists g \in q.\ (\mathrm{dom}(f) \cap \mathrm{dom}(g) = \emptyset)\ \wedge \\ f \sqsubseteq h_1\ \wedge\ \ldots\ \wedge\ f \sqsubseteq h_n\ \wedge\ g \sqsubseteq h_1\ \wedge\ \ldots\ \wedge\ g \sqsubseteq h_n.$$

Now, define $f_i = f$ and $g_i = h_i - f$ for $i \in \{1, \ldots, n\}$. Then,

$$(\forall i \in \{1, \ldots, n\}.\ f_i \cdot g_i = h_i) \\ \wedge\ (f_1, \ldots, f_n) \in \Delta_n(p)\ \wedge\ (g_1, \ldots, g_n) \in \Delta_n(q).$$

Hence, $(h_1, \ldots, h_n) \in \Delta_n(p) * \Delta_n(q)$. This shows that $\Delta_n(p * q) \subseteq \Delta_n(p) * \Delta_n(q)$. For the other inclusion, suppose that

$$(h_1, \ldots, h_n)\ \in\ \Delta_n(p) * \Delta_n(q).$$

Then, by the definition of $*$,

$$\exists (f_1, \ldots, f_n) \in \Delta_n(p).\ \exists (g_1, \ldots, g_n) \in \Delta_n(q). \\ (\forall i \in \{1, \ldots, n\}.\ f_i \cdot g_i = h_i).$$

Since $(f_1, \ldots, f_n) \in \Delta_n(p)$ and $(g_1, \ldots, g_n) \in \Delta_n(q)$, there are $f \in p$ and $g \in q$ such that

$$f \sqsubseteq f_1\ \wedge\ \ldots\ \wedge\ f \sqsubseteq f_n\ \wedge\ g \sqsubseteq g_1\ \wedge\ \ldots\ \wedge\ g \sqsubseteq g_n.$$

Furthermore, since f_1 and g_1 have disjoint domains, their subheaps f and g must have disjoint domains as well. Consequently, $f \cdot g$ is well defined, and it satisfies

$$f \cdot g \in p * q\ \wedge\ (\forall i \in \{1, \ldots, n\}.\ f \cdot g\ \sqsubseteq\ f_i \cdot g_i\ =\ h_i).$$

This implies that $(h_1, \ldots, h_n) \in \Delta_n(p * q)$, as desired. □

Lemma 3.1. *For all φ and η, ρ, ρ', if $\Delta_n(\rho(a)) = \rho'(a)$ for every $a \in \mathsf{AVar}$, we have that $\Delta_n([\![\varphi]\!]^1_{\eta,\rho}) = [\![\varphi]\!]^n_{\eta,\rho'}$.*

Proof. We prove by induction on the structure of φ. All the inductive cases and the cases of true and false follow from the preservation result of Lemma 2.2. Thus, it is sufficient to show the lemma when $\varphi \equiv a$ or $\varphi \equiv P$. When $\varphi \equiv a$, the assumption of the lemma implies that

$$\Delta_n([\![a]\!]^1_{\eta,\rho})\ =\ \Delta_n(\rho(a))\ =\ \rho'(a)\ =\ [\![a]\!]^n_{\eta,\rho'}.$$

When $\varphi \equiv P$, we note that $\Delta_n \circ \Delta_1 = \Delta_n$, and conclude that

$$\Delta_n([\![P]\!]^1_{\eta,\rho})\ =\ \Delta_n(\Delta_1((\!|P|\!)_\eta))\ =\ \Delta_n((\!|P|\!)_\eta)\ =\ [\![P]\!]^n_{\eta,\rho'}.$$

□

Appendix B. Proof of Lemma 4.1

Lemma 4.1. *There is an algorithm taking simple assertions φ, ψ and returning finitely many implications $\{\varphi^l \Longrightarrow \psi^l\}_{l\in L}$, such that (a) $\varphi^l \Longrightarrow \psi^l$ has the form (4.1) and (b) for any $n \in \{1, 2\}$, we have that $\varphi \models^n \psi$ holds iff $\varphi^l \models^n \psi^l$ holds for all $l \in L$.*

Proof. The algorithm first transforms ψ in the conjunctive normal form, using proof rules in classical logic, which hold in all the n-ary semantics. This gives an implication of the form:

$$\bigvee_{i=1}^{I} \bigwedge_{j=1}^{J} \varphi_{(i,j)} * \mathbf{a}_{(i,j)} \implies \bigwedge_{k=1}^{K} \bigvee_{l=1}^{L} \psi_{(k,l)} * \mathbf{b}_{(k,l)}.$$

Then, the algorithm constructs the below set:

$$\left\{ \bigwedge_{j=1}^{J} \varphi_{(i,j)} * \mathbf{a}_{(i,j)} \implies \bigvee_{l=1}^{L} \psi_{(k,l)} * \mathbf{b}_{(k,l)} \right\}_{1\le i\le I, 1\le k\le K}.$$

Finally, it removes, in each implication, all the disjuncts that include assertion variables not appearing on the LHS of the implication; if all disjuncts are removed, false is the new RHS. The outcome of this removal becomes the result of the algorithm. □

Appendix C. Layouts that Lift

Lemma C.1 (Segregation)**.** *For any $I, J \ge 1$ there are non-empty, finite* segregating *subsets $\mathbf{S}^{I,J}_{i,j} \subseteq \mathsf{PosInt}$ for all $1 \le i \le I$ and $1 \le j \le J$ with these properties:*

(1) $\forall 1 \le i_1, i_2 \le I.\ \ \bigcup_{1\le j\le J} \mathbf{S}^{I,J}_{i_1,j} = \bigcup_{1\le j\le J} \mathbf{S}^{I,J}_{i_2,j}$.
(2) $\forall 1 \le i \le I.\ \ \forall 1 \le j_1 \ne j_2 \le J.\ \ \mathbf{S}^{I,J}_{i,j_1} \cap \mathbf{S}^{I,J}_{i,j_2} = \emptyset$.
(3) $\forall 1 \le i_1 \ne i_2 \le I.\ \ \forall 1 \le j_1, j_2 \le J.\ \ \mathbf{S}^{I,J}_{i_1,j_1} \cap \mathbf{S}^{I,J}_{i_2,j_2} \ne \emptyset$.

By 1 we define $\mathbf{S}^{I,J} = \bigcup_{1\le j\le J} \mathbf{S}^{I,J}_{i,j}$ for any $1 \le i \le I$.

Theorem 4.5 (Shadow-Lift). *Unary η-validity of an implication implies the PC if each dashed line has a label that is not a label on a solid line and each disjunct has an occurrence of a variable that is not a label on a solid line. Spelling it out in symbols, we require, with $L = \{(i, j) \mid 1 \le i \le M \wedge 1 \le j \le N\}$, that*

$$\begin{array}{l} \forall (i,j) \in L.\ \ \Pi(i) \not\ge \Omega(j) \implies \\ \quad \exists c \in V.\ \ \Pi(i)(c) < \Omega(j)(c) \ \wedge \\ \qquad \big(\forall (k,l) \in L.\ \ \Pi(k) \ge \Omega(l) \implies \Pi(k)(c) = \Omega(l)(c)\big) \end{array}$$

and

$$\begin{array}{l} \forall 1 \le j \le N.\ \ \exists c \in V.\ \ \Omega(j)(c) > 0 \ \wedge \\ \quad \big(\forall (k,l) \in L.\ \ \Pi(k) \ge \Omega(l) \implies \Pi(k)(c) = \Omega(l)(c)\big). \end{array}$$

Proof. Assume that we have an implication of the form (4.1) in Section 4 and an appropriate environment η, that the stated criterion on the variable layout holds and that we have unary η-validity. We must show that the PC holds.

According to Definition 4.3 we assume that we have heaps $h, h_1, \ldots, h_M \in \mathsf{Heap}$ with $h_i \sqsubseteq h$ and $h_i \in [\![\varphi_i]\!]^1_\eta$ for all $1 \leq i \leq M$. The core of the proof is the construction of a particular environment $\rho : V \to \mathsf{IRel}_1$. For that purpose we need some notation. For a subset $M \subseteq \mathsf{PosInt}$ we denote by $[M]$ the heap that has domain M and stores some fixed value, say 0, at all these locations. Let $C \subseteq V$ be the set of assertion variables that do not occur as labels on solid edges, i.e., for a $c \in V$ we have that $c \in C$ iff

$$\forall 1 \leq i \leq M.\ \forall 1 \leq j \leq N.$$
$$\Pi(i) \geq \Omega(j) \implies \Pi(i)(c) = \Omega(j)(c).$$

For each $1 \leq i \leq M$ we let K_i be the set of second indices of all variables in conjunct i that lie in C, i.e., we set $K_i = \{1 \leq k \leq M_i \mid a_{i,k} \in C\}$. If non-empty, we let $k_i = \min(K_i)$.

We now define $\rho(c) = \mathsf{Heap}$ for $c \in V \setminus C$. For a variable $c \in C$ we let $\rho(c)$ be the union of

$$\bigcup_{\substack{1 \leq i \leq M,\\ K_i \neq \emptyset,\\ a_{i,k_i} \equiv c}} \left\{ (h - h_i) \cdot [\mathbf{S}^{M,K}_{i,k_i} + L] \cdot \prod_{\substack{1 \leq k \leq K,\\ k \notin K_i}} [\mathbf{S}^{M,K}_{i,k} + L] \right\}^{\uparrow}$$

and

$$\bigcup_{1 \leq i \leq M, K_i \neq \emptyset, k \in K_i \setminus \{k_i\}, a_{i,k} \equiv c} [\mathbf{S}^{M,K}_{i,k} + L],$$

where we have used $K = \max\{M_1, \ldots, M_M\}$ and $L = \max(\mathrm{dom}(h) \cup \{0\})$. For each $1 \leq i \leq M$ we can write $h \cdot [\mathbf{S}^{M,K} + L]$ as the following product

$$h_i \cdot (h - h_i) \cdot \prod_{k \in K_i} [\mathbf{S}^{M,K}_{i,k} + L] \cdot \prod_{1 \leq k \leq K, k \notin K_i} [\mathbf{S}^{M,K}_{i,k} + L],$$

which implies that we have $h \cdot [\mathbf{S}^{M,K} + L]$ a member of $[\![\varphi_i * a_{i,1} * \cdots * a_{i,M_i}]\!]^1_{\eta,\rho}$. In summary, we have shown that $h \cdot [\mathbf{S}^{M,K} + L]$ lies in the unary interpretation of the left hand side in the environments η and ρ. By assumption, the same must hold for the right hand side and from this we aim to derive the PC.

We now know that $h \cdot [\mathbf{S}^{M,K} + L]$ lies in the interpretation of some disjunct, say disjunct j. This means that

$$\begin{aligned} h \cdot [\mathbf{S}^{M,K} + L] &\in [\![\psi_j * b_{j,1} * \cdots * b_{j,N_j}]\!]^1_{\eta,\rho} \\ &= [\![\psi_j]\!]^1_\eta * \prod_{k \in J} \rho(b_{j,k}), \end{aligned}$$

where $J = \{1 \leq k \leq N_j \mid b_{j,k} \in C\}$ is the set of second indices of variables of disjunct j that are in C. By the second assumption of the theorem we know that $J \neq \emptyset$. We write

$$h \cdot [\mathbf{S}^{M,K} + L] = g \cdot \prod_{k \in J} g_k$$

for $g \in [\![\psi_j]\!]^1_\eta$ and $g_k \in \rho(b_{j,k})$ for each $k \in J$. By the properties of segregating sets we get that there must be a common $1 \leq i \leq M$ such that for all $k \in J$ there is $l_k \in K_i$ with

$$[\mathbf{S}^{M,K}_{i,l_k} + L] \sqsubseteq g_k,$$

i.e., the g_k's are all 'from the same conjunct'. But this implies $\Pi(i)(c) \geq \Omega(j)(c)$ for all $c \in C$ as the segregating sets are non-empty. But then $\Pi(i)(c) \geq \Omega(j)(c)$ must hold for $c \in V \setminus C$ too by the first assumption of the lemma and so $\Pi(i) \geq \Omega(j)$. Also we must have $\Pi(i)(c) = \Omega(j)(c)$ for each $c \in C$ by definition of C. By construction we have

$$\mathrm{dom}\left(\prod_{k\in J} g_k\right) \supseteq \mathrm{dom}(h - h_i) \cup (\mathbf{S}^{M,K} + L)$$

But then $\mathrm{dom}(g) \subseteq h_i$ and so we have $h_i \in [\![\psi_j]\!]^1_\eta$ too and we have proved the first option of the PC. □

Theorem 4.7 (Balloon-Lift). *Unary η-validity of an implication implies the PC if there is a subset $B \subseteq V$ with the following three properties. First, each conjunct has at most one occurrence of a variable from B, i.e.,*

$$\forall 1 \leq i \leq M.\ \sum_{c\in B} \Pi(i)(c) \leq 1.$$

Second, each disjunct is empty or has exactly one occurrence of a variable from B, i.e.,

$$\forall 1 \leq j \leq N.\ \sum_{c\in V} \Omega(j)(c) = 0 \ \vee \sum_{c\in B} \Omega(j)(c) = 1.$$

Third, each dashed line must have a label from B. That is, when $L = \{(i,j) \mid 1 \leq i \leq M \wedge 1 \leq j \leq N\}$,

$$\forall (i,j) \in L.\ \Pi(i) \not\geq \Omega(j) \implies \exists c \in B.\ \Pi(i)(c) < \Omega(j)(c).$$

Proof. Assume that we have an implication of the form (4.1) in Section 4 and an appropriate environment η, that the stated criterion on the variable layout holds and that we have unary η-validity. We must show that the PC holds.

According to Definition 4.3 we assume that we have heaps $h, h_1, \ldots, h_M \in \mathsf{Heap}$ with $h_i \sqsubseteq h$ and $h_i \in [\![\varphi_i]\!]^1_\eta$ for all $1 \leq i \leq M$. The core of the proof is the construction of a particular environment $\rho : V \to \mathsf{IRel}_1$. We define $\rho(c) = \mathsf{Heap}$ for $c \in V \setminus B$. For a variable $c \in B$ we let $\rho(c)$ be the following union

$$\bigcup_{1\leq i\leq M, 1\leq k\leq M_i, a_{i,k}\equiv c} \{h - h_i\}^\uparrow.$$

For each $1 \leq i \leq M$ we can write $h = h_i \cdot (h - h_i)$ and so we have h in $[\![\varphi_i * a_{i,1} * \cdots * a_{i,M_i}]\!]^1_{\eta,\rho}$ by the first of the original assumptions on the set B. In summary, we have shown that h lies in the unary interpretation of the left hand side in the environments η and ρ. By assumption, the same must hold for the right hand side and from this we aim to derive the PC.

We now know that h lies in the interpretation of some disjunct, say disjunct j. If this disjunct is empty we have proved the second option of the PC. Otherwise we know that there is exactly one $1 \leq k \leq N_j$ such that $b_{j,k} \in B$. But then we have

$$h \in [\![\psi_j * b_{j,1} * \cdots * b_{j,N_j}]\!]^1_{\eta,\rho} = [\![\psi_j]\!]^1_\eta * \rho(b_{j,k}).$$

We write

$$h = g \cdot g_k$$

for $g \in [\![\psi_j]\!]^1_\eta$ and $g_k \in \rho(b_{j,k})$. There must be an $1 \leq i \leq M$ such that $g_k \sqsupseteq h - h_i$ and such that $\Pi(i)(b_{j,k}) = \Omega(j)(b_{j,k}) = 1$. The first gives $h_i \in [\![\psi_j]\!]^1_\eta$ and the second implies $\Pi(i) \geq \Omega(j)$ by the third assumption on B. And we have arrived at the first option of the PC. □

Appendix D. Completeness

Lemma D.1 (Fan-Counter)**.** *Suppose that the layout of variables is as follows. There are at least two conjuncts, i.e., $M \geq 2$, and one conjunct has the property that each variable occurring in the conjunct also occurs as a label of a solid line leaving the conjunct and ending in a non-empty disjunct. In symbols the latter is*

$$\begin{aligned}\exists 1 \leq i \leq M.\ \forall c \in V.\ \Pi(i)(c) > 0 \implies \\ \exists 1 \leq j \leq N.\ \Pi(i) \geq \Omega(j) \wedge \Pi(i)(c) > \Omega(j)(c) \wedge \\ \exists d \in V.\ \Omega(j)(d) > 0.\end{aligned}$$

Then there are choices of φ_i's and ψ_j's with no variables such that the implication holds on the unary level but not on the binary level.

In the search for counterexamples we may without loss of generality assume the negation of the conditions of the above lemma. This means, provided at least two conjuncts, that for any non-empty set of solid lines leaving one common conjunct and ending in non-empty disjuncts there is a variable that occurs in the conjunct but is not a label of either of the lines. If, loosely phrased, we invalidate that variable, then all the solid lines break down, i.e., become dashed.

Lemma D.2 (X-Counter)**.** *Suppose that the layout of variables is as follows. There are two distinct conjuncts i_0 and i_1 and two distinct non-empty disjuncts j_0 and j_1 such that $\Pi(i_0) \not\geq \Omega(j_0)$ while $\Pi(i_0) \geq \Omega(j_1)$, $\Pi(i_1) \geq \Omega(j_0)$ and $\Pi(i_1) \geq \Omega(j_1)$. Then there are choices of φ_i's and ψ_j's that the implication holds on the unary level but not on the binary level.*

Again we may without loss of generality assume that the negation of this lemma holds when building counterexamples. Picture the graph of the implication without empty disjuncts and without dashed lines. The negation of the above means that we may arrive at all vertices in the connected component containing some vertex by paths from that vertex of length 2 or less. Also all connected components are complete, in particular no two vertices with a dashed line between them can belong to the same component.

Lemma D.3 (Bridge-Counter)**.** *Suppose that the layout of variables is as follows. There are at least two conjuncts, i.e., $M \geq 2$, all disjuncts are non-empty and there is a dashed line with labels that all occur as labels on solid lines too. In symbols the last demand is*

$$\begin{aligned}\exists 1 \leq i \leq M.\ \exists 1 \leq j \leq N.\ \Pi(i) \not\geq \Omega(j) \wedge \\ \forall c \in V.\ \Pi(i)(c) < \Omega(j)(c) \implies \\ \exists 1 \leq k \leq M.\ \exists 1 \leq l \leq N. \\ \Pi(k) \geq \Omega(l) \wedge \Pi(k)(c) > \Omega(l)(c).\end{aligned}$$

Then there are choices of φ_i's and ψ_j's with no variables such that the implication holds on the unary level but not on the binary level.

This lemma deals with the case of a variable layout with at least two conjuncts and no empty disjuncts but where the first condition of Theorem 4.5 fails.

Lemma D.4 (All-Out-Counter)**.** *Suppose that the layout of variables is as follows. There are at least two conjuncts, i.e., $M \geq 2$, at least one non-empty disjunct and for each variable one of the following two holds: Either the variable occurs as a label on a solid line ending in a non-empty disjunct. Or it occurs at least twice in a conjunct and we have an empty disjunct. In symbols the variable condition is*

$$\begin{aligned}\forall c \in V.\ \big(&\exists 1 \leq i \leq M.\ \exists 1 \leq j \leq N.\ \Pi(i) \geq \Omega(j) \wedge \\ &\Pi(i)(c) > \Omega(j)(c) \wedge \exists d \in V.\ \Omega(j)(d) > 0\big) \vee \\ &\big(\exists 1 \leq i \leq M.\ \Pi(i)(c) \geq 2 \wedge \\ &\quad \exists 1 \leq j \leq N.\ \forall d \in V.\ \Omega(j)(d) = 0\big).\end{aligned}$$

Then there are choices of φ_i's and ψ_j's φ_i's and ψ_j's with no variables such that the implication holds on the unary level but not on the binary level.

This lemma deals with two cases. The first is the case of a variable layout with at least two conjuncts and no empty disjuncts but where the second condition of Theorem 4.5 fails while the first holds. The second is the case of a variable layout with at least two conjuncts, at least one empty disjunct and no dashed lines for which Theorem 4.7 fails.

Appendix E. Higher Arities and Parametricity

Theorem 5.1. *For an implication of the form* (4.1) *and an appropriate environment η we have that n-ary η- validity implies the PC if $n \geq \max\{2, M_1, \ldots, M_M\}$.*

Proof. Assume that we have an implication of the form (4.1) in Section 4 and an appropriate environment η, that $n \geq \max\{2, M_1, \ldots, M_M\}$ and that we have n-ary η-validity. We must show that the PC holds.

According to Definition 4.3 we assume that we have heaps $h, h_1, \ldots, h_M \in \mathsf{Heap}$ with $h_i \sqsubseteq h$ and $h_i \in [\![\varphi_i]\!]^1_\eta$ for all $1 \leq i \leq M$. The core of the proof is the construction of a particular environment $\rho : V \to \mathsf{IRel}_n$. For that purpose we need some notation. Define, for each $1 \leq k \leq n$, a map $\gamma_k : \mathsf{Heap} \to \mathsf{Heap}^n$ by letting

$$\gamma_k(h) = (\overbrace{[], \ldots, []}^{k-1}, h, \overbrace{[], \ldots, []}^{n-k})$$

for any $h \in \mathsf{Heap}$, i.e., it returns the n-tuple that has h as the k-th entry and the empty heap everywhere else. Similarly, we define $\delta : \mathsf{Heap} \to \mathsf{Heap}^n$ by setting

$$\delta(h) = (\overbrace{h, \ldots, h}^{n})$$

for any $h \in \mathsf{Heap}$, i.e., it returns the n-tuple that has h as all entries. For a subset $M \subseteq \mathsf{PosInt}$ we denote by $[M]$ the heap that has domain M and stores some fixed value, say 0, at all these locations.

For a variable $c \in V$ we now define $\rho(c)$ to be the following union of relations in IRel_n:

$$\bigcup_{1 \leq i \leq M, 1 \leq k \leq M_i, a_{(i,k)} \equiv c} \{\gamma_k(h - h_i) \cdot \gamma_1([\mathbf{S}^{M,K}_{i,k} + L])\}^\uparrow$$

where $K = \max\{M_1, \ldots, M_M\}$, $L = \max(\mathrm{dom}(h))$. This is well-defined because of our assumption that $n \geq \max\{M_1, \ldots, M_M\}$. For each $1 \leq i \leq M$ we have that

$$\begin{aligned}\delta(h) &= \delta(h_i) \cdot \gamma_1(h - h_i) \cdot \cdots \cdot \gamma_n(h - h_i) \\ &\sqsupseteq \delta(h_i) \cdot \gamma_1(h - h_i) \cdot \cdots \cdot \gamma_{M_i}(h - h_i)\end{aligned}$$

where we use the extension order for heap tuples defined by pointwise extension in all entries. Also, we have that

$$\begin{aligned}[\mathbf{S}^{M,K} + L] &= [\mathbf{S}^{M,K}_{i,1} + L] \cdot \cdots \cdot [\mathbf{S}^{M,K}_{i,K} + L] \\ &\sqsupseteq [\mathbf{S}^{M,K}_{i,1} + L] \cdot \cdots \cdot [\mathbf{S}^{M,K}_{i,M_i} + L].\end{aligned}$$

This gives us that $\delta(h) \cdot \gamma_1([\mathbf{S}^{M,K} + L])$ extends the following n-tuple of heaps:

$$\delta(h_i) \cdot \prod_{1 \leq k \leq M_i} \gamma_k(h - h_i) \cdot \gamma_1([\mathbf{S}^{M,K}_{i,k} + L])$$

which again means that $\delta(h) \cdot \gamma_1([\mathbf{S}^{M,K} + L])$ lies in

$$[\![\varphi_i * a_{i,1} * \cdots * a_{i,M_i}]\!]^n_{\eta,\rho}.$$

In summary, we have shown that $\delta(h) \cdot \gamma_1([\mathbf{S}^{M,K} + L])$ lies in the n-ary interpretation of the left hand side in the environments η and ρ. By assumption, the same must hold for the right hand side and from this we aim to derive the PC.

There is $1 \leq j \leq N$ such that we have

$$\delta(h) \cdot \gamma_1([\mathbf{S}^{M,K} + L]) \in [\![\psi_j * b_{j,1} * \cdots * b_{j,N_j}]\!]^n_{\eta,\rho}.$$

Consider first the case of a non-empty disjunct, i.e., the case $N_j > 0$. We split along the disjunct and get

$$\delta(h) \cdot \gamma_1([\mathbf{S}^{M,K} + L]) = \delta(g) \cdot \mathbf{g}_1 \cdot \cdots \cdot \mathbf{g}_{N_j}$$

for $g \in [\![\psi_j]\!]^1_\eta$ and $\mathbf{g}_k \in \rho(b_{j,k})$ for all $1 \leq k \leq N_j$. By the properties of segregating sets we get that there must be a common $1 \leq i \leq M$ such that for all $1 \leq k \leq N_j$ there is $1 \leq k_k \leq M_i$ with

$$\gamma_{k_k}(h - h_i) \cdot \gamma_1([\mathbf{S}^{M,K}_{i,k_k} + L]) \sqsubseteq \mathbf{g}_k,$$

i.e., the $\mathbf{g}_k$'s are all 'from the same conjunct'. But this implies $\Pi(i) \geq \Omega(j)$ as the segregating sets are non-empty. Also the above equality enforces $\mathrm{dom}(g) \subseteq \mathrm{dom}(h)$ by the definition of γ_1. Indeed we must have $\mathrm{dom}(g) \subseteq \mathrm{dom}(h_i)$ since in particular we have

$$\gamma_{1_k}(h - h_i) \cdot \gamma_1([\mathbf{S}^{M,K}_{i,1_k} + L]) \sqsubseteq \mathbf{g}_1.$$

But then $g \sqsubseteq h_i$ so we have $h_i \in [\![\psi_j]\!]^1_\eta$ too and the first option of the PC holds.

We consider now the case of an empty disjunct, i.e., the case $N_j = 0$. As above we split along the disjunct and get

$$\delta(h) \cdot \gamma_1([\mathbf{S}^{M,K} + L]) = \delta(g) \cdot \mathbf{g}$$

for $g \in [\![\psi_j]\!]^1_\eta$ and $\mathbf{g} \in \mathsf{Heap}^n$. Again we must have $\mathrm{dom}(g) \subseteq \mathrm{dom}(h)$ which implies $g \sqsubseteq h$ and the second option of the PC holds. □

Appendix F. Proof of Theorem 6.1

Theorem 6.1. *Every derivable* $\Gamma \vdash \{\varphi\}C\{\psi\}$ *is* 2*-valid.*

Proof. We will show that all the rules in Figure 4 are sound. This lets us prove the theorem by induction on the height of the derivation of a judgment, because using the soundness of the rules, we can handle all the base and inductive cases.

Let's start with the rule for the module operation k. Suppose that we have $(\eta, \rho, \mathbf{u}) \models^2 (\Gamma, \{\varphi\}k\{\psi\})$. Then, by the definition of $\models^2$, we should have $(\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}k\{\psi\}$ as well. From this follows the soundness of the rule.

Next, consider four rules: (a) the frame rule for adding $- * \varphi$ to the pre and post-conditions, (b) the rule for adding $\exists x$ to the pre and post-conditions, (c) the rule for sequencing, and (d) the rule for the conditional statement. All these rules are sound, because of the following four facts:

$$\begin{array}{l}(\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}C\{\varphi'\} \\ \qquad \Longrightarrow\ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi * \psi\}C\{\varphi' * \psi\}\end{array}$$

$$\begin{array}{l}(x \notin \mathsf{FV}(C))\ \wedge\ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}C\{\psi\} \\ \qquad \Longrightarrow\ (\eta, \rho, \mathbf{u}) \models^2 \{\exists x.\varphi\}C\{\exists x.\psi\}\end{array}$$

$$\begin{array}{l}(\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}C\{\varphi'\}\ \wedge\ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi'\}C'\{\psi\} \\ \qquad \Longrightarrow\ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}C; C'\{\psi\}\end{array}$$

$$\begin{array}{l}(\eta, \rho, \mathbf{u}) \models^2 \{\varphi \wedge B\}C\{\psi\}\ \wedge \\ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi \wedge \neg B\}C'\{\psi\} \\ \qquad \Longrightarrow\ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi\}\mathbf{if}\ B\ C\ C'\{\psi\}\end{array}$$

The first fact is an easy consequence of using the quantification over IRel_2 in the semantics of triples. The second also follows easily from the semantics of triples and the fact that $[\![C]\!]_{\eta,u_i} = [\![C]\!]_{\eta[x\mapsto v],u_i}$ for all $v \in \mathsf{Int}$, as long as one remembers that the $*$ operator distributes over union. The third and fourth are not different, and they follow from the semantics of triples and commands. Here we will go through the details of proving the fourth fact. Consider $(\eta, \rho, \mathbf{u})$ satisfying the assumption of the fact. Pick $r \in \mathsf{IRel}_2$ and heaps f, g such that

$$(f, g)\ \in\ [\![\varphi]\!]^2_{\eta,\rho} * r.$$

Now we do the case analysis on whether $[\![B]\!]_\eta$ is true or not. If it is true, then

$$\begin{array}{rl}(f, g)\ \in & [\![\varphi \wedge B]\!]^2_{\eta,\rho} * r \\ \wedge & [\![\mathbf{if}\ B\ C\ C']\!]_{\eta,u_1}(f)\ =\ [\![C]\!]_{\eta,u_1}(f) \\ \wedge & [\![\mathbf{if}\ B\ C\ C']\!]_{\eta,u_2}(g)\ =\ [\![C]\!]_{\eta,u_2}(g).\end{array}$$

Hence, by assumption, we get that

$$\begin{array}{c}\big([\![\mathbf{if}\ B\ C\ C']\!]_{\eta,u_1}(f),\ [\![\mathbf{if}\ B\ C\ C']\!]_{\eta,u_2}(g)\big) \\ =\ \big([\![C]\!]_{\eta,u_1}(f),\ [\![C]\!]_{\eta,u_2}(g)\big)\ \in\ [\![\psi]\!]^2_{\eta,\rho} * r.\end{array}$$

If $[\![B]\!]_\eta$ is not true, we reason similarly, but with C' instead of C, and get that

$$\big([\![\mathbf{if}\ B\ C\ C']\!]_{\eta,u_1}(f),\ [\![\mathbf{if}\ B\ C\ C']\!]_{\eta,u_2}(g)\big)\ \in\ [\![\psi]\!]^2_{\eta,\rho} * r.$$

We have just shown that in both cases, the outcomes of the conditional statements are related by $[\![\psi]\!]^2_{\eta,\rho} * r$, as claimed by the fourth fact.

We move on to the rules for heap update and dereference. They are sound because of the below two facts:

$$(\eta, \rho, \mathbf{u}) \models^2 \{E \hookrightarrow _\}[E] := F \{E \hookrightarrow F\}$$

$$\begin{array}{l} x \notin \mathsf{FV}(\psi) \ \ \wedge \ \ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi * E \hookrightarrow x\} C \{\psi\} \\ \quad \Longrightarrow \ (\eta, \rho, \mathbf{u}) \models^2 \{\exists x. \varphi * E \hookrightarrow x\} \mathbf{let}\ x = [E]\ \mathbf{in}\ C \{\psi\} \end{array}$$

To prove the first, we pick $(\eta, \rho, \mathbf{u})$, a relation $r \in \mathsf{IRel}_2$ and heaps f, g such that

$$(f, g) \ \in \ [\![E \hookrightarrow _]\!]^2_{\eta,\rho} * r.$$

Then, there exist heaps h, f_1, g_1 such that

$$(f, g) = (h, h) \cdot (f_1, g_1) \ \ \wedge \ \ [\![E]\!]_\eta \in \mathrm{dom}(h) \ \ \wedge \ \ (f_1, g_1) \in r.$$

Thus,

$$\begin{array}{l} (\, [\![[E] := F]\!]_{\eta,u_1}(f), \ \ [\![[E] := F]\!]_{\eta,u_2}(g) \,) \\ \quad = \ \big(h[[\![E]\!]_\eta \mapsto [\![F]\!]_\eta] \cdot f_1, \ \ h[[\![E]\!]_\eta \mapsto [\![F]\!]_\eta] \cdot g_1 \big) \\ \quad \in \ [\![E \hookrightarrow F]\!]^2_{\eta,\rho} \ * \ r, \end{array}$$

as desired by the first fact. For the proof of the second fact, suppose that the assumption of the second fact holds, and pick $r \in \mathsf{IRel}_2$ and heaps f, g such that

$$(f, g) \ \in \ [\![\exists x. \, \varphi * E \hookrightarrow x]\!]^2_{\eta,\rho} * r.$$

Then, there exists an integer v and heaps h, f_1, g_1, f_2, g_2 such that

$$\begin{array}{l} (f, g) = (h, h) \cdot (f_1, g_1) \cdot (f_2, g_2) \ \ \wedge \ \ h \in [\![E \hookrightarrow x]\!]^1_{\eta[x \mapsto v]} \\ \qquad \wedge \ \ (f_1, g_1) \in [\![\varphi]\!]^2_{\eta[x \mapsto v],\rho} \ \ \wedge \ \ (f_2, g_2) \in r \end{array}$$

Thus, $(f, g) \in [\![\varphi * E \hookrightarrow x]\!]^2_{\eta[x \mapsto v],\rho} \ * \ r$, and $f([\![E]\!]_\eta) = g([\![E]\!]_\eta) = v$. Using these and the assumed triple of the second fact, we derive the below:

$$\begin{array}{l} (\, [\![\mathbf{let}\ x = [E]\ \mathbf{in}\ C]\!]_{\eta,u_1}(f), \ \ [\![\mathbf{let}\ x = [E]\ \mathbf{in}\ C]\!]_{\eta,u_2}(g) \,) \\ \quad = \ \big([\![C]\!]_{\eta[x \mapsto v],u_1}(f), \ \ [\![C]\!]_{\eta[x \mapsto v],u_2}(g) \big) \\ \quad \in \ [\![\psi]\!]^2_{\eta[x \mapsto v],\rho} * r \\ \quad = \ [\![\psi]\!]^2_{\eta,\rho} * r. \end{array}$$

The last equality holds, because x does not appear in φ. We have just proved that the output states of two dereferencing commands are $([\![\psi]\!]^2_{\eta,\rho} * r)$-related, as claimed by the second fact.

Finally, we prove that the rule of consequence is sound. It is sufficient to show that

$$\begin{array}{l} \textsc{Chk}(\varphi', \varphi) \ \ \wedge \ \ \varphi' \models^1 \varphi \ \ \wedge \\ \textsc{Chk}(\psi, \psi') \ \ \wedge \ \ \psi \models^1 \psi' \ \ \wedge \ \ (\eta, \rho, \mathbf{u}) \models^2 \{\varphi\} C \{\psi\} \\ \quad \Longrightarrow \quad (\eta, \rho, \mathbf{u}) \models^2 \{\varphi'\} C \{\psi'\}. \end{array}$$

From the first four conjuncts of the assumption, it follows that

$$\varphi' \models^2 \varphi \quad \wedge \quad \psi \models^2 \psi'.$$

This is due to the correctness of CHK, which holds because all the transformations used in the first check of CHK are based on semantic equivalences holding in IRel_2 and the second lifting check is sound because of our lifting theorems. In order to prove the conclusion of the above implication, pick $r \in \mathsf{IRel}_2$ and heaps f, g such that

$$(f, g) \ \in \ [\![\varphi']\!]^2_{\eta,\rho} * r.$$

Since the $*$ operator is monotone and $\varphi' \models^2 \varphi$, we get that

$$(f, g) \ \in \ [\![\varphi]\!]^2_{\eta,\rho} * r.$$

This relationship and the assumed triple $\{\varphi\}C\{\psi\}$, then, imply the below:

$$(\,[\![C]\!]_{\eta,u_1}(f),\, [\![C]\!]_{\eta,u_2}(g)\,) \ \in \ [\![\psi]\!]^2_{\eta,\rho} * r.$$

Again, since $\psi \models^2 \psi'$, the monotonicity of the $*$ operator implies that

$$(\,[\![C]\!]_{\eta,u_1}(f),\, [\![C]\!]_{\eta,u_2}(g)\,) \ \in \ [\![\psi']\!]^2_{\eta,\rho} * r.$$

Note that this is the conclusion that we are looking for. □

This work is licensed under the Creative Commons Attribution-NoDerivs License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/2.0/ or send a letter to Creative Commons, 171 Second St, Suite 300, San Francisco, CA 94105, USA, or Eisenacher Strasse 2, 10777 Berlin, Germany